GridPP Log Book – QCDgrid quarterly report, Sep. 2007 (Q3/07)

Preface

This report provides a status update on project progress at end of Project Month 37 (Sept. 2007). It is based upon the updated project work plan that was approved by GridPP in October 2006, and includes proposed additions to the work plan to extend the project duration until the end of January 2008.

1. Project

Name: QCDgrid

Manager: Professor Richard Kenway, The University of Edinburgh.

Status of milestones:

4.4	**UKQCD**			**Status Date 30-Sep-07**
Owner:	Richard Kenway			
Number	**Title**	**Due**	**Status**	**Links**
4.4.1	WP1.4 - First QCDGrid software review	01-Mar-05	Complete	WP 1.4
4.4.2	WP1.4 - Second QCDGrid software review	01-Jul-06	Complete	WP 1.4
4.4.3	WP1.6 - Test, extend and maintain the current QCDgrid software	On-going	OK – 7 releases	WP 1.6
4.4.4	WP2.4 - Incorporate QCDML marked-up data from QCDOC production codes.	01-Sep-05	Complete	WP 2.4
4.4.5	WP3.4 - Design of metadata catalogue Web Service	01-Dec-05	Complete	WP 3.4
4.4.6	WP1.7 - Training for UKQCD users *Added in Jul. '05 Change Request*	01-May-05	Complete	WP 1.7
4.4.7	WP4.2 – Functional Specification of ILDG File Catalogue service *Changed in Dec. '05 Change Request*	01-Jan-07	Complete	WP 4.2
4.4.8	WP4.7 - Dissemination and x-coordination with ILDG community *Deleted in Oct. '06 Change Request*	01-Jul-07	Deleted	
4.4.9	WP5.1 - Web service incorporating both the metadata and replica catalogues. *Re-scheduled in Oct '06 Change Request*	01-Jul-07	Complete	WP 5.1
4.4.10	WP 6.4 - Implementation of revisions to QCDgrid Suite. *Added in Oct '06 Change Request*	01-Feb-08	In progress	WP 6.4

Change Request History

Date	Summary of Change Request
July 2005	Replacement of WP 3.6 and WP 3.7 with new WP 1.7 (and new milestone 4.4.6).
December 2005	Refinement of WP 3 and WP 4 work packages.
October 2006	Extension of project and creation of new work package WP 6.

2. High Level Objectives and Level 1 Deliverables

Objective I

(a) **Description:** To improve, maintain and develop a simple and easy mechanism for UKQCD scientists to share and access each other's data and compute resources.

(b) **Purpose:** Quantum Chromodynamics (QCD) is an application area that requires access to the most powerful supercomputing resources and generates large amounts of raw data, called configurations, which require extensive post-processing. QCD scientists require secure and reliable access, from geographically dispersed locations, to these computer systems and data.

(c) **Principal client:** UKQCD scientists and GridPP (QCDgrid source code is freely available).

(d) **Definition of successful achievement of objective:** UKQCD scientists continue to be able to generate, access and process data in a seamless way across a secure, reliable and extensible grid from a range of geographically dispersed locations. This will involve monitoring of technology developments, both in terms of hardware on the grid and in terms of software developments in the underlying grid technologies. In addition, metadata will be automatically generated alongside configuration generation.

(e) **High level risks:** R2, R7, R10, R38. Lack of continuing support for underlying grid technologies, particularly the Globus toolkit and QCDgrid software. A detailed risk analysis will be carried out in Work Package 1.

(f) **Level 1 Deliverables:**

a. End of PM 33 – Test, extend and maintain the current QCDgrid software throughout the duration of the project to, for example, cope with rapid data generation when QCDOC begins production. This will involve:

   i. a review of requirements and stress testing of the software by end of PM 5;

   ii. two code reviews of current technologies such as EDG/EGEE and the Globus toolkit, to be completed by end of PM 6 and end of PM 22 (Project Map deliverables 4.4.1 and 4.4.2);

   iii. an on-going maintenance and upgrade effort (Project Map deliverable 4.4.3).

b. End of PM 10 – incorporate QCDML marked-up data from QCDOC production codes (Project Map deliverable 4.4.4).

c. End of PM 8 – organisation and delivery of a UKQCD HackLatt workshop, producing training on the usage and administration of UKQCD resources, with a particular focus on the UKQCD Grid and QCDOC (Project Map deliverable 4.4.6).

Objective II

(a) **Description:** To provide a simple and easy mechanism for QCD scientists across different continents to share and access each other's data.

(b) **Purpose:** Quantum Chromodynamics (QCD) is an application area that requires access to the most powerful supercomputing resources and generates large amounts of raw data, called configurations, which require extensive post-processing. QCD scientists require to share this data across different continents in a secure and reliable way.

(c) **Principal Client:** World-wide lattice QCD community via the International Lattice Data Grid (ILDG).

(d) **Definition of successful completion of objective:** International members of the ILDG collaboration are able to securely and reliably access data from geographically dispersed locations. An easy to access replica catalogue exists across multiple countries. An easy to access metadata catalogue exists across multiple countries.

(e) **High level risks:** R2, R3, R7, R35. Failure to achieve deliverables in Objective I. Delays in agreeing standards amongst ILDG members. Lack of effort in other countries to implement data grid software.

(f) **Level 1 Deliverables:**

a. End of PM 22 – Produce web service implementation of the metadata catalogue, built on top of the QCDgrid Metadata Catalogue, including web interface (Project Map deliverable 4.4.5).

b. End of PM 32 – Demonstrate web service implementation of the replica catalogue, built on top of the QCDgrid Replica Catalogue, including web interface (Project Map deliverable 4.4.7 and 4.4.8).

c. End of PM 33 – Demonstrate web service-based grid solution incorporating both the metadata catalogue and replica catalogue across UK and non-UK nodes (Project Map deliverable 4.4.9).

3. Level 2 Deliverables or Milestones

WP 1 Level 1: Test, extend and maintain the current QCDgrid software throughout the duration of the project

Effort:	18.3 months
Duration:	41 months (Due end of PM 41 – January 2008)
Inputs:	See Level 2 tasks
Related req's:	Input from the UKQCD collaboration
Products:	See Level 2 deliverables
Status	IN PROGRESS (WP1.5 and WP1.6 are active, at time of writing).

This work package will extend, test, and maintain the current QCDgrid software throughout the duration of the project to, for example, cope with rapid data generation when QCDOC begins production.

WP 1.1 Risk analysis

Effort:	1 month
Duration:	2 months (Due end of PM 2 – October 2004)
Inputs:	
Related req's:	Input from UKQCD collaboration
Products:	Risk analysis document
Status	COMPLETE – Work Package 1.1: QCDgrid System Risk Analysis published on GridPP website in October 2004.

This deliverable will carry out a risk analysis to identify potential problems and failures of the QCDgrid software. The analysis will consider probably changes in its use over the duration of the project, covering (for example): acquisition of new hardware, changes to usage patterns, and new internationally-agreed standards.

WP 1.2 Stress testing requirement capture

Effort:	1 month
Duration:	2 months (Due end of PM 3 – November 2004)
Inputs:	WP 1.1
Related req's:	Input from UKQCD collaboration
Products:	Test Suite requirements definition document
Status	COMPLETE – Work Package 1.2: QCDgrid Stress Testing Requirements report published on GridPP website in November 2004.

This deliverable consists of a formal requirements capture of the UKQCD collaboration's needs for a QCDgrid Software Stress Test Suite.

WP 1.3 Stress testing

Effort:	1 month
Duration:	3 months (Due end of PM 5 – January 2005)
Inputs:	WP 1.2
Related req's:	-
Products:	Implementation of Software Test Suite
Status	COMPLETE: • QCDgrid Stress Test Suite available from NeSCForge in CVS module 'testing'. • Work Package 1.3: QCDgrid Stress Testing Results report published on GridPP website (February 2005).

This deliverable consists of the implementation and documentation of stress tests to determine the behaviour of the QCDgrid software and the UKQCD Grid under high demand conditions, as identified during the requirements capture work package WP 1.2.

WP 1.4 Software review (PM4.4.1 and PM4.4.2)

Effort:	2 months
Duration:	4 months (2 months due end of PM 6 – Feb. 2005, and 2 months due end of PM 24 – Aug. 2006)
Inputs:	
Related req's:	Input from UKQCD collaboration
Products:	Software review document ×2
Status	COMPLETE – Work Package 1.4.1 Software Review (Part 1) report published on GridPP website (March 2005). COMPLETE – Work Package 1.4.2 Software Review (Part 2) report published on GridPP website (September 2006).

Two software reviews of the underlying QCDgrid software have been carried out at appropriate points in the project timeline. These have consider software such as the EGEE application stack and newer version of the Globus Toolkit. The conclusions drawn from these two reviews feed into subsequent releases of the QCDgrid software. To aid this, metrics will be established and documented within the review that allow us to monitor the successful adoption of conclusions within QCDgrid software.

WP 1.5 Maintenance and support

Effort:	9.8 months
Duration:	41 months (Due end of PM 41 – Jan. 2008)
Inputs:	
Related req's:	Input from UKQCD collaboration
Products:	A reliable and secure QCDgrid and UKQCD contribution to ILDG throughout the project. Quarterly reports will be provided on this activity – see below for details.
Status	IN PROGRESS – see description, below, for details of specific activities.

This deliverable encapsulates the on-going maintenance and support of the QCDgrid software throughout the duration of the project. This includes the base QCDgrid software suite along with additional resources and software that are generated for the UKQCD contribution to the ILDG.

This deliverable has fulfilled the following specific activities[1]:

- **January 2005**: Establishment of a bug/query/support request tracking environment, hosted by NeSCForge (see http://forge.nesc.ac.uk/projects/qcdgrid/).

[1] This list of activities is being populated with notable items, as the Level 2 deliverable progresses.

- **March 2005**: Completion of WP1.5 *QCDgrid Software Release Plan* that documents a framework for providing new releases of the QCDgrid software suite with four month frequency.
- **April 2005**: Successful completion of UKQCD HackLatt 2005 workshop, which provided a unique opportunity for the QCDgrid team to:
    - disseminate training on the usage and administration of the QCDgrid applications;
    - interact with and gather feedback from the UKQCD user community.

  The following material is available from the workshop:
    - UKQCD HackLatt 2005 Final Report;
    - Workshop presentation: The QCDgrid software suite, an introduction;
    - Workshop presentation: QCDgrid Administration.
- **July 2005**: During 05Q3, technical problems were encountered with the UKQCD Grid that necessitated higher than expected levels of effort expenditure on this work package. More details of these problems can be found in the corresponding subsection of Section 4.
- **February 2006**: Deliverable D1.5.2 UKQCD Grid Support Review 2005 was released. This document provides a review of the support-related activities of the project, during the period from September 2004 to the end of December 2005. It examines significant events and issues, proposing modifications to the support process in order to improve future performance and streamline the effort expenditure of the QCDgrid project.
- **March 2006**: Successfully completed QCDgrid System Administration Training 2006, that brought together system administrators and software developers from across the UK academic community and beyond, with the intention of providing training on the deployment, maintenance and development of QCDgrid-powered services. A number of documents are available from the workshop webpage, including the final report and copies of workshop presentations.
- **November 2006**: To reduce diagnostic time associated with failures of grid services and resources, the project team have deployed the GITS monitoring software onto the UKQCD Grid. This provides a status page covering availability of basic grid services (GridFTP, GRAM, and RLS) and the amount of free space on UKQCD storage resources.
- **January 2007**: The XML database used to host the UKQCD metadata catalogue has been upgraded from eXist Version 0.9 to eXist Version 1.1. The newer release is no longer classed as a beta release, implying improved robustness and reliability. Version 1.1 also introduces a number of performance enhancements that help to reduce query time – especially when the database is accessed via the ILDG metadata catalogue web service.
- **May 2007**: The EDG VO LDAP service, which had been used to centralise the administration of UKQCD user information, has been replaced by a VOMS service. In addition to UKQCD user information, the VOMS service hosts user information for the whole of the ILDG community. The migration (from LDAP to VOMS) has been completed in stages: it began in January 2007 and was completed in May 2007. Local software to support the VOMS system has been developed by the project team within WP 4.1.

WP 1.6 Software Installation and Distribution (PM4.4.3)

Effort:	2 months
Duration:	41 months (Due end of PM 41 – Jan. 2008)
Inputs:	WP 1.3, WP1.4.
Related req's:	Input from UKQCD collaboration
Products:	Roll out of packaged, easy-to-install updates to QCDgrid software, throughout the lifetime of the project. The following updates have thus far been produced: • QCDgrid Version 1.2, released during May 2005. • QCDgrid Version 1.3, released during October 2005. • QCDgrid Version 1.4, released May 2006. • QCDgrid Version 1.5, released September 2006. • ILDG Browser 1.6, released March 2007. • DiGS Version 2.0, released October 2007.
Status	IN PROGRESS – see description, below, for details of specific activities.

During 2007Q3, the team released DiGS Version 2.0. This is a major update to the QCDgrid software, which includes numerous new and extended functionalities (see the announcement on NeSCForge for more details). The most significant contribution of DiGS Version 2.0, from the perspective of the project objectives, is that it enables access to data for our ILDG collaborators.

WP 1.7 UKQCD HackLatt Workshop (PM 4.4.6)

Effort:	1.5 months
Duration:	2 months (Due end of PM 8 – Apr. 2005)
Inputs:	Contributions from WP 1.5 and WP 1.6.
Related req's:	Input/effort from UKQCD collaboration
Products:	This work package will generate two products: • A two day workshop in April 2005 (PM 8), delivering training on the usage and administration of UKQCD computing resources, with a particular focus on the UKQCD Grid and QCDOC. • UKQCD HackLatt 2005 Final Report (May 2005).
Status	COMPLETE – workshop information is available from the workshop webpage.

The project team organised a workshop entitled *UKQCD HackLatt 2005*. This workshop brought together particle physicists and software developers from across the UK and beyond, with the aim

of disseminating information and training on the powerful and diverse computational resources that are available within the UK, focusing on UKQCD Grid and QCDOC.

WP 2 Level 1: Incorporate QCDML marked-up data from QCDOC production codes.

Effort:	3.6 months
Duration:	5 months (Due end of PM 12 – Aug. 2005)
Inputs:	
Related req's:	Input from UKQCD collaboration and ILDG, QCDML XML application
Products:	See Level 2 deliverables.
Status	COMPLETE, less WP2.3 which has been SUSPENDED.

Web-based (or other such) tools will be developed to support and simplify the usage of QCDML XML application by scientists.

WP 2.1 Metadata tools requirements capture

Effort:	1 month
Duration:	3 months (Due end of PM – Jun. 2005)
Inputs:	
Related req's:	Input from UKQCD collaboration and ILDG, QCDML XML application
Products:	Requirements Definition document.
Status	COMPLETE – Work Package 2.1: QCDgrid Metadata Tools Requirements Capture, report published on GridPP website (June 2005).

This deliverable records the results of a requirement capture exercise carried out for the work package – focusing on metadata creation and management tools. The exercise was conducted by the QCDgrid project team at EPCC, in conjunction with members of the UKQCD collaboration. The outputs of the requirements capture exercise were as follows:

1. a full understand of the structure of the metadata used by UKQCD scientists;
2. an understand of the processes required to create and manipulate the metadata, during the day-to-day activities of QCD scientists;
3. an identification of the gaps in functionality within the existing software that prevents the key processes identified in item 2 from being carried out effectively/efficiently;
4. A definition of the requirements, in terms of new software and enhancements to existing software, to bridge the gaps identified in item 3.

WP 2.2 Tooling to support QCDML metadata management

Effort:	1 month
Duration:	2 months (Due end of PM 11 – Jul. 2005)
Inputs:	WP 2.1
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	New release of QCDgrid Metadata Catalogue Browser (or, possibly, a separate tool) that incorporates functionality to allow the submission of ensemble/gauge configuration metadata to the QCDgrid MDC.
Status	COMPLETE -- revised application source code is available from NeSCForge via CVS.

This work package has designed and implemented revised versions of the QCDgrid clients (both command-line and GUI-based). These tools support Use Cases 1a and 2a from the requirement identified in WP 2.1. They will be formally released as part of the next major version, QCDgrid Version 1.3, publication of which is imminent.

WP2.2 expended 2 months of effort – as opposed to the planned one month effort. Because of this, effort available to other components of the work package has been reduced: see discussion of WP2.3 below for details.

WP 2.3 Tooling to support QCDML metadata creation

Effort:	1 month
Duration:	2 months (Due end of PM 12 – Aug. 2005)
Inputs:	WP 2.1
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	Web-based (or other such) tooling to support the creation of ensemble metadata documents and gauge configuration metadata documents.
Status	SUSPENDED

The level of work involved in WP 2.2 was higher than anticipated, due to a more extensive and complex implementation phase. As a consequence, assigned effort on Work Package 2 has now been exhausted. The project team have recommended that WP2.3 be suspended, in order to: keep the project on track with respect to the work plan; allow sufficient time to work on WP3.1 in preparation for the ILDG Middleware workshop in Tsukuba at the end of October 2005.

Adverse effects caused by the suspension of WP2.3 can be mitigated by the deployment and use of third-party XML editing tools, such as XML Spy and the Eclipse XML plug-in. The project team will monitor any problems that arise as a result of using generic application, instead of bespoke, in-house tools and respond accordingly.

WP 2.4 Dissemination and liaison with ILDG (PM4.4.4)

Effort:	0.6 month
Duration:	5 months (Due end of PM 12 – Aug. 2005)
Inputs:	WP 2.2, WP 2.3
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	User documentation for WP 2.2 and WP 2.3.
Status	COMPLETE – revised user documentation is available from NeSCForge via CVS. This documentation will be incorporated in the next release version of the QCDgrid (Version 1.3) which is expected during October 2005.

The aim of this work package is to publicise and disseminate information about the revisions to the datagrid client tools that have been made during WP2.1, WP2.2, and WP2.3, targeting both the UKQCD user community and the ILDG.

With regard to the UKQCD user community, revised user documentation has been prepared to complement the revised tools. This documentation will feed into the next formal version of the QCDgrid software.

In addition, the project team have disseminated information about our work to the ILDG through a meeting with Jim Simone (from Fermi National Accelerator Laboratory in Chicago, USA) on Wednesday 3rd August 2005. During the meeting, the team discussed the status and future efforts of the ILDG activities, plus considered how the QCDgrid client could be adopted as a standard client tool across the ILDG. This discussion will be pursued at the next ILDG Middleware workshop in Tsukuba, Japan at the end of October 2005.

WP 3 Level 1: Web service implementation of ILDG Metadata Catalogue Interface

Effort:	8 months
Duration:	21 months (Due end of PM 22 – Jun. 2006)
Inputs:	
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	Demonstration of implementation of ILDG web service interface to the QCDgrid Metadata Catalogue, supported with appropriate documentation.
Status	COMPLETE.

In this work package, the QCDgrid project team will liaise with the ILDG community to define a web service interface for the ILDG-wide Metadata Catalogue (MDC). Having achieved agreement on the definition of the interface, the project team will design and implement an appropriate extension to the QCDgrid Metadata Catalogue, to fulfil the ILDG specification and enable metadata created by UKQCD to be published via the ILDG infrastructure.

WP 3.1 Contributions to ILDG Middleware Working Group workshops

Effort:	1.5 month (3×0.5 months)
Duration:	18 months (Due end of PM 20 – Apr. 2006)
Inputs:	Contributions from UKQCD collaboration.
Related req's:	
Products:	• Minutes from ILDG Middleware Workshop (Edinburgh, Scotland, 25th—27th October 2004). Available as a PDF document from the ILDG website. • Minutes from ILDG Middleware Workshop (Tsukuba, Japan, 27th—28th October 2005). Available as a PDF document from the ILDG website. • Minutes from ILDG Middleware Workshop (JLab, USA, 10th—13th December 2006). Available as a PDF document (temporary link, as ILDG website is down at time of writing).
Status	COMPLETE

Specification of the interface(s) to be supported by the ILDG MDC (and File Catalogue) is (are) within the remit of the ILDG Metadata Working Group. The project has established a representation on this working group in the form of George Beckett and Daragh Byrne.

The working group typically meets face-to-face, once every twelve months. Three meetings of the working group have been convened, at which the ILDG MDC (and File Catalogue) have been defined. These are:

- Workshop #1 in Edinburgh, UK (October 2004);
- Workshop #2 in Tsukuba, Japan (October 2005);
- Workshop #3, in Virginia, USA (December 2006).

WP 3.1 encapsulates the effort required to contribute to these three meetings (expected to be approximately 0.5 months/meeting). The primary output from each workshop is a comprehensive minutes document, which is published on the ILDG website. These minutes and the agreements that they describe form the basis of subsequent work in WP 3 (and WP 4).

WP 3.2 Establishment of development environment for work package

Effort:	0.5 month
Duration:	1 months (Due end of PM 15 – Nov. 2005)
Inputs:	Workstation with development specification (to be provided by EPCC).
Related req's:	Minutes from Workshop #1 and #2; QCDgrid Test Grid (as established in WP 1.3).
Products:	Revised QCDgrid Test Grid infrastructure, ready to support definition, implementation and testing for WP 3.5.
Status	COMPLETE.

The purpose of WP 3.2 is to establish an appropriate project environment for the development tasks required by subsequent packages within WP 3 (and WP 4). The environment has been built on top of the existing test bed infrastructure – from hereon referred to as the *QCDgrid Development System* – that is in place at EPCC. The following notable modifications have been made:

- Addition of a fourth node to the QCDgrid Development System, with an appropriate specification for web service development.
- Installation of the Eclipse development environment.
- Installation and testing of the Tomcat Apache Servlet container and Axis SOAP engine.
- Installation of latest version of XML database used to power UKQCD Grid MDC (that is, Apache Exist).

WP 3.3 Specification of web service interface to the MDC (Version 1)

Effort:	1 month
Duration:	2 month (Due end of PM 15 – Nov. 2005)
Inputs:	
Related req's:	Minutes from Workshop #1 and #2
Products:	D3.3: Specification of the web service interface to the MDC (Version 1) – report published on project web pages (November 2005).
Status	COMPLETE.

The discussions and agreements of the ILDG Middleware working group have been refined into a detailed specification of the ILDG MDC web service interface (Version 1) to be implemented. Amongst other things, this specification provides:

- a definition of the functionality for the first version of the web service interface;
- a set of use cases that cover the functionality of the ILDG MDC;
- acceptance criteria for a valid web service interface for the ILDG MDC.

WP 3.4 Design and architecture for QCDgrid implementation of ILDG MDC Version 1 (PM4.4.5)

Effort:	1 month
Duration:	2 month (Due end of PM 17 – Jan. 2006)
Inputs:	
Related req's:	D3.3: Specification of web service interface to MDC (Version 1)
Products:	D3.4: Design for QCDgrid implementation of ILDG MDC (Version 1) (Jan. 2006)
Status	COMPLETE.

In WP 3.4, the project team have completed a design proposal for the extension to the QCDgrid package, required to fulfil the specification of the ILDG MDC web service interface (Version 1).

The work package also provides an appraisal of deployment requirements for integrating the UKQCD Grid (MDC) into the ILDG service.

WP 3.5 Implementation of ILDG-compliant web service layer (Version 1) in QCDgrid

Effort:	1 month
Duration:	1 month (Due end of PM 18 – Feb. 2006)
Inputs:	
Related req's:	D3.4.1: Design for QCDgrid implementation of ILDG MDC (Version 1)
Products:	ILDG-compliant web service implementation for the UKQCD Grid.
Status	COMPLETE.

An implementation of the ILDG MDC web service interface has been developed in WP 3.5. This is currently deployed onto the QCDgrid Development System and is being tested by our ILDG Collaborators.

Contrary to the original proposal, it has been determined that the web service interface should not be packaged with the core QCDgrid software but instead distributed as an add-in. This decision has been made in order to maintain the general applicability of the QCDgrid software – the web service interface is a domain-specific component highly tailored to the needs of the ILDG.

WP 3.6 Collation of feedback from user community on ILDG MDC (Version 1)

Effort:	0.5 month
Duration:	1 month (Due end of PM 19 – Mar. 2006)
Inputs:	Input from ILDG and UKQCD user community.
Related req's:	QCDgrid Version 1.4
Products:	D3.6: Feedback from Community on ILDG Metadata Catalogue Web Service Version 1 (May 2006)
Status	COMPLETE

In Work Package 3.6, the project team have appraised the status of the ILDG MDC Web Service Version 1, as the ILDG members approach the realisation of a production service. The appraisal includes:

- A chronology of the key events in the development of the service.
- A status report on the situation, at the time of writing.
- A number of proposals for continued development and production deployment of the service across the ILDG community.

One of the key conclusions of the report relates to the development of a fully functioning web service client that is tailored to a typical user. We are pleased to report that ILDG have agreed that the QCDgrid Browser is an appropriate client, though we need to make modifications to the back-end connectivity to integrate it into the ILDG service. To this end, the project team have replaced WP3.7, WP3.8 and WP3.9, as previously defined, with a specification of work for the ILDG client.

WP 3.7 Design of web service client to the ILDG MDC (Version 1)

Effort:	1 month
Duration:	2 month (Due end of PM 21 – May 2006)
Inputs:	Contributions from ILDG Middleware Working Group Existing QCDgrid Browser client source code.
Related req's:	D3.6: Feedback from Community on ILDG Metadata Catalogue Web Service Version 1
Products:	Work Package 3.7: Metadata Catalogue Web Service Browser Design (June 2006)
Status	COMPLETE, revised version released in August 2006.

This work package determines the process of modifying the QCDgrid GUI Browser in order that it can be deployed as the principal client interface to the ILDG Metadata Catalogue.

This involves an investigation of the following specific areas:

- The current features of the existing browser that are desirable in an ILDG version of the browser;

- The current features of the existing browser that are undesirable/unnecessary in an ILDG version of the browser;
- How existing features need to be modified to interact with web services instead of XML databases;
- How undesirable/unnecessary existing features could be disabled in the ILDG version of the browser

The design establishes a factoring of source code that permits parallel versions of the client to co-exist for UKQCD (the full QCDgrid Browser) and ILDG (lightweight implementation with read-only functionalities) from a common source code base.

WP 3.8 Implementation of web service client to the ILDG MDC (Version 1)

Effort:	1 month
Duration:	1 month (Due end of PM 22 – Jul. 2006)
Inputs:	
Related req's:	Work Package 3.7: Metadata Catalogue Web Service Browser Design
Products:	ILDG Metadata Web Service client, available for download from NeSCForge (Aug. 2006).
Status	COMPLETE.

WP 3.8 encapsulates the implementation of the ILDG Metadata Catalogue Web Service Client, a stand-alone Java client that works on Solaris/Linux systems that have a JVM (Version 1.4 or higher). It has also been demonstrated to work with Microsoft Windows, though this O/S is not officially supported by the team.

WP 4 Level 1: Web service implementation of ILDG File Catalogue

Effort:	9 months
Duration:	11 months (Due end of PM 35 – Jul. 2007)
Inputs:	See Level 2 tasks.
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	Demonstration of Web service implementation of the File Catalogue, including web-based (or otherwise) client application. User documentation.
Status	COMPLETE.

This work package involves specification, design and development of a web service interface to the UKQCD Grid File Catalogue (FC) and Storage Elements (SEs), built on top of the existing QCDgrid software.

WP 4.1 Migration of UKQCD VO to ILDG VO

Effort:	1 month
Duration:	2 months (Due end of PM 26 – Oct. 2006)
Inputs:	
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	ILDG VO service (hosted at https://grid-voms.desy.de:8443/voms/ildg/). ILDG VO Policy Statement document (August 2006 – co-authored by members of ILDG Middleware Working Group). Revised version of VO authentication within the QCDgrid software (available from 'tools' module of project's NeSCForge software repository).
Status	COMPLETE.

In conjunction with our ILDG collaborators, we have establish an ILDG-wide Virtual Organisation, which is hosted by DESY Zeuthen (German Grid) and supported by the ILDG VO Policy Statement. The ILDG VO has been implemented using the EGEE VOMS system. The QCDgrid security infrastructure has been modified to source authorisation information from the ILDG VO, using a modified version of the EGEE VOMS command-line client tools.

WP 4.2 Functional specification of ILDG File Catalogue service (PM 4.4.7)

Effort:	1 month
Duration:	2 months (Due end of PM 28 – Dec. 2006)
Inputs:	(Experiences gained from) D3.3: Specification of the web service interface to the MDC (Version 1) and D3.6: Feedback from Community on ILDG Metadata Catalogue Web Service Version 1.
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	WP4.2 Functional Specification of ILDG File Catalogue (Version 1.0, January 2007). ILDG File Catalogue Web Service Operations (Version 0.1, July 2006).
Status	COMPLETE.

In conjunction with our ILDG collaborators, we have defined and documented a detailed specification of the ILDG File Catalogue web service interface to be implemented. The file catalogue is a core component of the ILDG infrastructure that is used to track the location of Lattice Gauge Configuration datasets within the storage resources provided by the regional grids.

WP 4.3 Security review of web service access to UKQCD Grid File Catalogue and Storage Elements

Effort:	1 month
Duration:	1 month (Due end of PM 29 – Jan. 2007)
Inputs:	WP 4.1, WP 4.2
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	WP4.3 ILDG File Catalogue Security (Version 1.0, February 2007).
Status	COMPLETE.

A security review has been conducted to ensure that the ILDG VO implementation and functional specification of the ILDG File Catalogue can be satisfied within the constraints of UKQCD data access policies and built on top of the existing QCDgrid software.

WP 4.4 Design and architecture for QCDgrid implementation of ILDG interface to File Catalogue and Storage Elements

Effort:	2 months
Duration:	2 months (Due end of PM 31 – Mar. 2007)
Inputs:	WP 4.1, WP 4.2, WP 4.3
Related req's:	-
Products:	WP4.4: Design and architecture for the QCDgrid implementation of the ILDG File Catalogue and Storage Element (May 2007).
Status	COMPLETE.

A design has been prepared for the ILDG File Catalogue, a key component of the ILDG infrastructure. The work focuses on the modifications that are required for the UKQCD Grid and associated software, to make the system compatible with the ILDG FC web service interface. The design has been broken down into three sections:

- security consideration for ILDG user authentication and authorisation.
- changes to the QCDgrid software, which are to be folded into a major new release called DiGS.
- web services required to realise the ILDG File Catalogue interface.

WP 4.5 Implementation of ILDG File Catalogue and basic Storage Element access

Effort:	2 months
Duration:	2 months (Due end of PM 33 – May 2007)
Inputs:	WP 4.1, WP 4.3, WP 4.4
Related req's:	
Products:	ILDG-compliant web service implementation for the UKQCD Grid.
Status	COMPLETE.

In this work package, we have implemented the design established in WP 4.3 and WP 4.4. The implementation is encapsulated into two software contributions:

- DiGS Version 2.0 (plus updated User Guide);
- ILDG File Catalogue Web Service.

Both of these are available from the project account on NeSCForge.

WP 4.6 Revised design for ILDG Browser to support File Catalogue interface

Effort:	1 month
Duration:	1 month (Due end of PM 34 – Jun. 2007)
Inputs:	WP 3.8, WP 4.5
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	ILDG FC Interface Service (0.3), prepared by ILDG Middleware Working Group (June 2007).
Status	COMPLETE.

The QCDgrid team – in conjunction with ILDG collaborators – have designed a command-line client for the ILDG File Catalogue web service (developed in WP 4.5). This is a change to the original implementation plan, since the command-line client will not integrate with the ILDG Browser client. A command-line client has been determined to be more appropriate, in the short term, since it supports both interactive and script-based usage.

WP 4.7 Implementation of revised ILDG Browser

Effort:	1 month
Duration:	1 month (Due end of PM 35 – Jul. 2007)
Inputs:	WP 4.5, WP 4.6
Related req's:	-
Products:	Command-line client for ILDG File Catalogue Web Service, available (via anonymous access) from ILDG CVS repository on NeSCForge (in module named "FC-WS-Client".[2]
Status	COMPLETE.

The QCDgrid team, in conjunction with ILDG collaborators – specifically, Dirk Pleiter and David Melkumyan – have implemented a command-line client for the ILDG File Catalogue web service (developed in WP 4.5). As noted in WP 4.6, this is a change to the original implementation plan, since the command-line client will not integrate with the ILDG Browser client. A command-line client has been determined to be more appropriate, in the short term, since it supports both interactive and script-based usage.

[2] Please note that, at the time of writing, software that is developed within the remit of the ILDG Middleware Working Group is only available via anonymous access to ILDG CVS repository on NeSCForge. The normal source (the ILDG website) has been suspended, by the hosting institution (Jefferson Laboratories, Virginia, USA) due to security concerns.

WP 5 Level 1: Demonstrate web service based grid solution incorporating both the metadata catalogue and file catalogue across UK and non-UK nodes

Effort:	1.1 months
Duration:	1 month (Due end of PM 35 – Jul. 2007)
Inputs:	See Level 2 deliverables
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	See Level 2 deliverables
Status	COMPLETE

WP 5.1 Final web based data grid solution (PM4.4.9)

Effort:	1.1 months
Duration:	1 month (Due end of PM 35 – Jul. 2007)
Inputs:	WP 2, WP 3, WP 4.
Related req's:	Input required from UKQCD collaboration and ILDG
Products:	• Demonstration at "The 16th International Symposium on Lattice Field Theory" in Regensburg, Germany (July 2007), of web service-based data-grid solution across multiple continents (including both metadata catalogue and file catalogue). • Conference paper, Towards an interoperable International Lattice Datagrid, prepared by the ILDG Middleware Working Group (inc. M.G. Beckett and R.H. Ostrowski) (July 2007)
Status	COMPLETE

This deliverable involves coordination with other ILDG groups to complete a demonstration of ILDG infrastructure at "The 16th International Symposium on Lattice Field Theory" in Regensburg, Germany (July 2007).

- A live demonstration of the ILDG infrastructure was successfully made, with connections to five different regional grids: specifically, UKQCD, CSSM (Australia), JLDG (Japan), LatFor Data Grid (Gernamy), and USQCD (USA).
- The demonstration has been documented in a paper, which has been submitted to the conference proceedings (see Products above).

WP 6 Level 1: Review of Storage Element deployment on UKQCD Grid

Effort:	4 months
Duration:	4 months (Due end of PM 41 – Jan. 2008)
Inputs:	See Level 2 tasks.
Related req's:	Input required from UKQCD collaboration and ILDG.
Products:	See Level 2 deliverables.
Status	IN PROGRESS, WP 6.1 and WP 6.2 are active.

This work package will conduct an appraisal of the existing UKQCD Grid Storage Elements, in light of the experiences of other applications within the GridPP programme and other collaborations within the ILDG.

WP 6.1 Review of Storage Element implementations

Effort:	1 month
Duration:	1 month (Due end of PM 36 – Aug. 2007)
Inputs:	WP 1.4 and WP 4.4
Related req's:	Input required from other applications in GridPP programme and ILDG.
Products:	Review of options for storage provision in UKQCD Grid
Status	IN PROGRESS.

Within this task, we will conduct an appraisal of the storage provision that has been undertaken by our co-collaborators in ILDG and compatible applications within the GridPP programme. We anticipate that our attention will focus on:

- Globus Reliable File Transfer service.
- SRM-based solutions, considering both disk-only and tape-based resources.

This is as described in D1.4.2.

WP 6.2 Migration plan for UKQCD Grid Storage Elements

Effort:	1 month
Duration:	1 month (Due end of PM 37 – Sep. 2007)
Inputs:	WP 6.1
Related req's:	Input required from UKQCD collaboration.
Products:	Migration plan for storage provision within UKQCD Grid
Status	IN PROGRESS.

Based on the output of WP 6.1, we will devise a migration plan for UKQCD Grid Storage Elements that aims to address the on-going data management requirements of UKQCD within the constraints of the software and hardware resources available.

WP 6.3 Revised design for QCDgrid suite to support revised Storage Element requirements

Effort:	2 months
Duration:	2 months (Due end of PM 39 – Nov. 2007)
Inputs:	WP 6.1, WP 6.2
Related req's:	
Products:	Revised design document for QCDgrid
Status	TO START.

This work package determines the modifications to the QCDgrid suite that are required to support the revised Storage Element requirements of UKQCD.

WP 6.4 Implementation of revisions to QCDgrid Suite (PM4.4.10)

Effort:	2 month
Duration:	2 months (Due end of PM 41 – Jan. 2008)
Inputs:	WP 6.3
Related req's:	-
Products:	Release of QCDgrid software that supports UKQCD requirements for Storage Elements.
Status	TO START.

This task encapsulates the revision of the QCDgrid software to support UKQCD requirements for Storage Elements.

4. Commentary

04Q3 Comments

Project progress

The second phase of the UKQCD work area (under the GridPP2 project family) began on 1st September 2004. Following the work plan described earlier in this document, we focused effort on two level 2 deliverables. Specifically, WP 1.1 and 0. These are both progressing according to expectation and more detailed information can be found under the corresponding level 2 descriptions.

The project also contributed to **two** conferences/meetings – details are provided in Section 5.

ILDG-related dissemination activities

The UKQCD collaboration have begun discussions with physicists in Orsay (France) who have recently joined the ILDG. It is intended that Olivier Brand-Foissac (Orsay) will visit Edinburgh during either October or November 2004 to learn about the UKQCD Grid infrastructure and options for adoption of a similar system for the French node of the ILDG data repository.

04Q4 Comments

Project progress

In line with the project work plan, the following level 2 deliverables have been completed and published to the project's webpage on GridPP: WP1.1 and WP1.2. Furthermore, good progress has been made with the following work packages:

- WP1.3 – a design for stress testing the QCDgrid software has been finalised, a local test grid (consisting of four nodes) has been assembled at EPCC, and implementation of stress testing is progressing.
- WP1.5/WP1.6 – a software release plan has been devised. Furthermore, preparations have been put in place to ensure that QCDgrid software can be installed on the QCDOC machine as soon as possible.

ILDG-related dissemination activities

Between 20th and 21st October 2004, EPCC hosted a **short visit by Olivier Brand-Foissac** (of the Université de Paris Sud in Orsay, France), acting as a representative for the new French component of the ILDG collaboration.

The purpose of Olivier's visit was to learn more about the activities of the UKQCD and ILDG. It was also a good opportunity to talk with the development team of the QCDgrid software. The visit was very successful and we anticipate continued communications into the future. It also looks likely that the French ILDG nodes will be powered by QCDgrid software – the collaboration is currently completing an acquisition of suitable hardware to support their data grid and associated services. Olivier provided a short report on his visit that is reproduced in Figure 1.

Date: Wed, 03 Nov 2004 19:13:50 +0100

From: Olivier Brand-Foissac <Olivier.Brand-Foissac@th.u-psud.fr>

To: Beckett George <george@epcc.ed.ac.uk>

Cc: LEROY_Jean-Pierre <Jean-Pierre.Leroy@th.u-psud.fr>,

BRAND-FOISSAC_Olivier <Olivier.Brand-Foissac@th.u-psud.fr>

Subject: EPCC visit

In order to setting up a french node into ILDG, a visit to EPCC at Edinburgh University was organized by George Beckett.

After a short introduction, the EPCC QCD Grid team (George Beckett, James Perry, Daragh Byrne, etc.) showed how UKQCD Grid uses QCD Grid software to store and retreive lattice QCD data from different sites in UK.

Based upon VDT packages with Globus-2 and a metadata data-base in XML (Exist / Tomcat), the QCD Grid software enables the users to search the metadata catalog from a Java interface and download the data files from the nearest grid site.

Tracks of files, consistency checking and replica location strategy were explained by James and Daragh together with the message passing across the grid.

Securized by certificates, the access of the grid functions propose a working panel of what could be a data grid based on Globus-2.

Chris Johnson's presentation of the ENACS experience using QCD Grid software as a middleware showed the versatility of the software beyond the difficulty of implementation on some host systems.

Richard Kenway explained the aims of ILDG and the links between QCD grids around the world, technically completed by Balint Joo with some of the development perspectives with the community.

The QCD Grid software, developped and supported at EPCC, already available on internet, documented and running in the UKQCD Grid, is a serious candidate for setting up a french data grid for QCD, joining soon the ILDG community.

Figure 1: Transcript of short report from Olivier Brand-Foissac, regarding his visit to EPCC in October 2004.

Between 25th—27th October 2004, Daragh Byrne and George Beckett attended the **ILDG Middleware Workshop** at NeSC. Details of the meeting can be found at http://www.nesc.ac.uk/esi/events/472/.

05Q1 Comments

Project progress

From the project work plan, the following level 2 deliverables have been completed during the first quarter of 2005:

- WP1.3 – a testing suite has been developed that allows the project team to simulate the UKQCD Grid under high demand, even to breaking point. The activity has confirmed that the QCDgrid software is reliable and predictable under realistic, high stress. Several issues have been raised by the activity, as follows:
    - scalability of QCDgrid Datagrid when multiple users attempt to write data to it is sub-optimal;
    - the eXist database used to host the QCDML metadata exhibits reliability issues under high load.

  Both of these issues will be monitored and corrective action will be taken over the next few months.

- WP1.4.1 – an audit of the QCDgrid software components has been conducted. The conclusion of this activity is that there are no immediate issues that affect our usage of third-party components within QCDgrid software. However, in the longer term, the project team aim to upgrade or replace these components. Specifically:
    - The eXist database (currently at Version 0.9) should be upgraded to Version 1.0 – when this version is released. Having upgraded the database, the reliability issues identified in WP1.3 will be revisited. If they persist, then the project team will evaluate alternative XMLDB API-compliant databases, such as Xindice.
    - The Grid middleware that QCDgrid is built upon (Globus Version 2.4) will cease to be supported when Globus Version 4.2 is released (timescale for this is not confirmed). Therefore, the QCDgrid team will undertake a thorough investigation of the options for upgrade/replacement of this middleware environment. Two candidates have been identified as most suitable: these are Globus Version 4.x and gLite Version 1.0. Neither of these packages is available at the time of writing. The project team will therefore conduct an evaluation of the two options once early adoption experiences from other projects have been produced. It should be noted that either approach to upgrading our middleware base will likely require significant effort.
    - It is likely that a richer Access Control Mechanism will be required in order to integrate the UKQCD Grid into ILDG. Currently, the requirements for such a mechanism are to be defined by the appropriate ILDG working groups. A thorough investigation of an appropriate mechanism is scheduled for Work Package 3 of the project. However, based on the early evaluation conducted in this work package, the Virtual Organisation Management System (VOMS) that is part of gLite, appears the most suitable candidate.
- WP1.5 – now that QCDgrid has reached a stable and complete implementation, it is appropriate to define a formal release strategy, designed to streamline future development and support. To this end, a WP1.5 QCDgrid Release Plan document has been prepared that defines two processes:
    - Process 1:– Software Versioning describing how updates to the QCDgrid software application suite are to be published;
    - Process 2:– problem resolution describing how problems with the QCDgrid software and issues with the UKQCD Grid infrastructure are to be tracked/assigned/resolved.

  In line with this plan, the project team intend to release Version 1.2 of QCDgrid during April 2005 – see Section 0 for details.

Note: WP1.2, WP1.3 and WP1.4.1 are tracked as Level 1 contributions to Objective I of the project.

Dissemination activities

The project will make two presentations at the up-coming UKQCD HackLatt 2005 conference, which is co-supported by GridPP and NeSC. These presentations will by targeted at QCDgrid users and administrators, respectively. Furthermore, Dr George Beckett is a co-organiser of this workshop.

During the first quarter of the year, the project team have held meetings with Professor Bill Allcock (GridFTP technical director, Argonne National Laboratories) and Professor Lisa Childers (Technical Production Manager for the Globus Alliance, Argonne National Laboratories). At these two meetings, QCDgrid made presentations of our Grid middleware usage and gathered information on the future plans for the Globus Toolkit, plus feedback on the options for upgrading our current

Grid middleware to the up-coming Globus Version 4.0. The output from these meetings has been documented in WP1.4.1.

05Q2 Comments

Project progress

During this quarter, the following progress has been made:

- WP1.5 – As part of the team's commitment to promote and support the usage of the QCDgrid software, a UKQCD HackLatt workshop was organised and delivered in April 2005. This workshop included a significant component allocated to the usage and administration of QCDgrid, plus instruction and advice on Grid computing in general. The workshop was co-supported by GridPP and NeSC. The project team made a number of contributions to the organisation and delivery of the workshop. These contributions are catalogued in section 0 on page 6. Following on from the workshop, the project team prepared a press release for GridPP that has subsequently been used for various GridPP publicity activities. The text of the press release is reproduced in Figure 2.

UKQCD gets to grips with the Grid at HackLatt 2005

Researchers from the UKQCD collaboration got to know the Grid better at the HackLatt workshop in Edinburgh, earlier this year. The workshop highlighted two key computing developments in the QCD field: the installation of the QCDOC supercomuter at Edinburgh, and UKQCD Grid (http://www.gridpp.ac.uk/qcdgrid). 27 scientists attended, including post-graduate students, researchers and lecturers from the fields of Physics and High Performance Computing.

Quantum chromodynamics (QCD) is the study of the building blocks of our universe. In the UK, and around the world, scientists are developing techniques to quantify the complex behaviour of fundamental particles called quarks and gluons - the constituents of all nuclear matter.

The UKQCD Collaboration (http://www.ph.ed.ac.uk/ukqcd) has been established to procure and develop computing resources to support numerical simulation of quantum chromodynamics. As part of this activity, the QCDOC super computer was installed at the University of Edinburgh in early 2005. The other key tool developed by the collaboration is UKQCD Grid, a multi-terabyte storage facility over six UK sites at Edinburgh, Liverpool, RAL, Southampton, and Swansea. Glasgow are also a member of the consortium.

HackLatt ran on 13th--14th April 2005: it was funded by GridPP and NeSC, and hosted by EPCC at the University of Edinburgh. The two day event provided a mixture of keynote presentations, hands-on tutorials, and discussion sessions. Key topics included:

- Getting to grips with the QCDOC super computer.
- Using QCD simulation codes on HPC resources.
- Exploiting the UKQCD Grid to manage and manipulate simulation datasets.

Physics and software engineering experts, including Professor Robert Edwards of the Jefferson Laboratories in Virginia, were on hand to provide advice and answer questions from participants.

Full details, including copies of presentations and tutorials, can be found on the UKQCD Collaboration website http://www.ph.ed.ac.uk/ukqcd. Based on the success of the event, UKQCD aim to provide follow-on events in years to come.

Figure 2: GridPP news article on UKQCD HackLatt 2005 workshop.

- WP1.6 – Version 1.2 of the QCDgrid software suite was released during May 2005. This is the first release of the project's software under the new release plan (as documented in WP1.5 QCDgrid Software Release Plan), which is designed to formalise a roadmap for the on-going maintenance of the software and simplify diagnostics and support of the UKQCD Grid infrastructure. Rollout of QCDgrid Version 1.2 across the UKQCD Grid servers and clients is continuing and, at the time of writing, we expect the rollout to be concluded by mid-July 2005.

- WP2.1 – a requirements capture exercise was executed and documented, covering the provision of tooling to support the creation and management of QCDML documents within the QCDgrid application software. The outputs from this exercise are documented in Work Package 2.1: QCDgrid Metadata Tools Requirements Capture.
- GridPP web pages – the structure of the project GridPP web pages has been revised to reduce the volume of material presented on the front page and, hopefully, make it simpler for a visitor to navigate through the project resources.

Dissemination activities

- Two presentations were prepared and delivered at the UKQCD HackLatt 2005 workshop, covering usage and administration of the QCDgrid software, as follows:
  - The QCDgrid software suite, an introduction
  - QCDgrid Administration.
- Revisions were made to the paper "QCDgrid: A Grid Resource for Quantum Chromodynamics" (submitted to Journal of Grid Computing) in response to referees comments. The revised text has now been returned to the editor for approval.
- Chris Maynard gave a presentation to the PPARC UK e-Science Postgraduate School, entitled "QCDgrid, a Grid for UKQCD". The presentation was included in a half-day session on Grid applications. See Section 5.2 for details of how to obtain a copy of the presentation.

05Q3 Comments

During this quarter, the project experienced some technical problems with the live UKQCD Grid that impacted on the amount of progress made on Work Package 2. The source of the technical problems was two-fold:

- **Problem 1:** Heavy user load on the UKQCD Grid control node (not related to QCDgrid software) caused a number of ambiguous failures of the system – these occurred when available memory on the control node dropped to critically low levels. This problem was not related to the QCDgrid software, though proved difficult to track down and expended significant support effort from the QCDgrid team.
- **Problem 2:** An inefficiency in the QCDgrid software was exposed by a dramatic increase in the number of datasets held on the UKQCD Grid – during the quarter approximately 5,000 new gauge configurations have been added. The inefficiency was in a verification code used to ensure the consistency of the Replica Catalogue, which scaled as $O(n^2)$ with the number of files on the grid. The offending algorithm is being replaced by a more efficient O(n log(n)) algorithm.

Between them, these two problems required 0.8 months of effort (instead of the predicted 0.3 months of effort), delaying work on Work Package 2 by approximately 0.5 months.

Despite this, good progress was made with WP 2.2 and WP2.4 – both of which are now complete. A new version of the QCDgrid client tools has been designed and implemented. This new version fulfils Use Cases 1a and 2a documented in Work Package 2.1: QCDgrid Metadata Tools Requirements Capture. It is available along with user documentation from NeSCForge via CVS. Furthermore, it will be incorporated into the QCDgrid Version 1.3 distribution expected during October 2005.

The level of work involved in WP 2.2 was higher than anticipated, due to a more extensive and complex implementation phase. As a consequence, assigned effort on Work Package 2 has now been exhausted. The project team have recommended that WP2.3 be suspended, in order to: keep

the project on track with respect to the work plan; allow sufficient time to work on WP3.1 in preparation for the ILDG Middleware workshop in Tsukuba at the end of October 2005.

Adverse effects caused by the suspension of WP2.3 can be mitigated by the deployment and use of third-party XML editor, such as XML Spy and the Eclipse XML plug-in. The project team will monitor any problems that arise as a result of using generic application, instead of bespoke, in-house tools and respond accordingly.

Looking to the future, the project team began to prepare for Work Package 3 and, in relation to this activity, George Beckett and Daragh Byrne made an application to the UK HEP Travel Grant to support their attendance at the International Lattice Data Grid Middleware Group Workshop 2005 at the University of Tsukuba, Tsukuba, Japan on 26th—27th October 2005.

Dissemination activities

- The project team are pleased to report that the paper "QCDgrid: A Grid Resource for Quantum Chromodynamics" (submitted to Journal of Grid Computing) has now been accepted for publication.
- The project team held a meeting with Jim Simone (from Fermi National Accelerator Laboratory in Chicago, USA) on Wednesday 3rd August 2005. The meeting discussed current progress and future plans of ILDG activities, plus identified areas of commonality between activities of this project and corresponding activities from other ILDG partners. These discussions will be pursued more fully at the up-coming workshop in Tsukuba and a detailed description will be provided in the next quarterly report.

05Q4 Comments

Project progress

During the final quarter of 2005, the project made good progress in line with the work plan. A new version (Version 1.3) of the QCDgrid software was released in October. This version incorporates new functionalities added during Work Package 2, along with several bug fixes identified in the Version 1.2 distribution. There were, however, two problems with the UKQCD Grid infrastructure that required a larger than planned amount of effort to be spent on support:

- A hardware failure on the Control Node led to a two week period of downtime. This was considered highly unsatisfactory and has led us to procure a new system to act as control node in the future. The new system will be a dedicated server, hosted within EPCC and covered by a next day maintenance contract. It is anticipated that the new system will be in operation by the end of February 2006.
- A number of issues have been encountered when attempting to add the University of Edinburgh Advance Computing Facility as a site on the UKQCD Grid. The target server node, which is called QCDOCx and is linked directly to the QCDOC's storage facility, has been the source of a number of problems, pertaining to software compatibility and network configuration. These problems are on-going, though the project team are liaising with QCDOC Systems staff to resolve the issues. We hope that the QCDOCx node will be working correctly by the end of January 2006.

ILDG-related activities

Daragh Byrne and George Beckett attended the ILDG Middleware Workshop, held in Tsukuba, Japan on 27th and 28th October 2005. The meeting proved very successful and, as a result, brought the ILDG close to agreement on the specification of the ILDG Metadata Catalogue service, as described in the Work Package 3 deliverable D3.3: Specification of the web service interface to the MDC (Version 1). The project team are now completing a design specification for the UKQCD

implementation of the interface and are optimistic that a pilot ILDG Metadata Catalogue will be in operation during Summer 2006. As part of the preparations for this, the project has completed the upgraded of the UKQCD Test Grid in line with the plans outlined in WP 3.2.

Dissemination activities

- The project team attended and contributed (via the ILDG Middleware Working Group) to the 7th ILDG Meeting, held in December 2005.
- In November, the UKQCD consortium participated in an international review of research using HPC in the UK. The QCDgrid project was a significant component of the UKQCD consortium contribution. The final report[3] included very positive comments about the UKQCD consortium and the project, including the following:

  > *"During site visits to UK institutions, the panel found that most of the research groups are aware of the need to focus on data, but do not have specific plans for or a comprehensive understanding of, how to handle data flexibly and efficiently. The UKQCD Consortium are the most advanced and articulated their data needs most clearly. They have initiated the International Lattice Data Grid (ILDG) and have developed Grid tools for semantic-based data access using metadata catalogues and XML schemata. UKQCD already have a middleware layer running, have successfully tested a prototype of their setup, and are now ready to go into production mode. [UKTC]"* – quoted from Section 4.3 Data Resources.
  >
  > *"Specific funding and a proposal mechanism should be developed to create a workforce that develops common tools and infrastructures for data handling, organisation, manipulation, and creation of user interfaces. The workforce should evaluate how much of the technology developed by the UKQCD consortium can be taken over for other consortia. The workforce should create a library of software for scientists dealing with large amounts of data."* – quoted from Section 5 Conclusions and Recommendations.

- The QCDgrid project team met with Dr Ann Chervenak (University of Southern California) in November 2005. Ann is the technical co-ordinator for the Globus Toolkit Replica Location Service. The project team described the functionality of the UKQCD Grid service and gathered important information on the future of the Replica Location Service component of the Globus Toolkit. This information will be particularly valuable as the project establishes plans for a Grid Middleware upgrade.

06Q1 Comments

Project progress (including ILDG-related activities)

During the first quarter of 2006, the project team has made good progress with Work Package 3 – designing (WP 3.4) and deploying (WP 3.5) a local implementation of the ILDG Metadata Catalogue web service (Version 1). The project team are also pleased to report significant progress made by a number of our ILDG collaborators. Specifically, both the LATFOR collaboration (DESY, Germany) and the Japanese collaboration (CCS, Tsukuba, Japan) have also deployed local implementations of the Version 1 web service.

---

[3] Horst Simon (Chair of International Review Panel), *International Review of Research Using HPC in the UK*, EPSRC and Deutsche Forschungsgemeinschaft (December 2005)

The QCDgrid project team have been co-ordinating interoperability tests between the three sites. The results from this testing were due for release at the end of March 2006, as output to WP 3.6. However, the results have been delayed by two weeks to: allow our ILDG collaborators a little more time to complete their implementation work; and to accommodate our contributions to the QCDgrid System Administration Training workshop (see below for details).

The QCDgrid project team are optimistic about the progress that has been made towards a functioning ILDG Metadata Catalogue web service. The team also hope to be able to share their web service implementation with USQCD, who – like UKQCD – have chosen the Apache eXist database to host their metadata, but who have not been able to progress their web service implementation during the previous quarter.

The project team also released project deliverable D1.5.2 UKQCD Grid Support Review 2005. The report provides a review of the support-related activities of the project, during the period from September 2004 to the end of December 2005. It examines significant events and issues, proposing modifications to the support process in order to improve future performance and streamline the effort expenditure of the QCDgrid project.

One recommendation of D1.5.2 is to deploy a dedicated server to act as control node for the UKQCD Grid, hosted at EPCC – the base location for the project development staff. Such a node was purchased in January 2006. At the time of writing, it has been configured appropriately for its new role and is awaiting deployment onto the UKQCD Grid.

Dissemination activities

- The QCDgrid project team are pleased to announce that the HPC-Europa NA3 project are currently evaluating the QCDgrid software suite as a candidate solution to their data management requirements.

  The project includes the following partners: Parallab, Bergen, Norway (leading activity); EPCC, Edinburgh, UK; TCD, Dublin, Ireland; CINECA, Bologna, Italy (leading HPC-Europa project); SARA, Amsterdam, Netherlands; and CASPUR, Rome, Italy.

  The general objective of Network Activity NA3 is to build competence and knowledge on existing tools and standards and in addition contribute to emerging tools and standards for efficient and portable data management and information retrieval in real-life computational science applications in global networks.

  We were pleased to welcome three members of the project development team to attend the QCDgrid System Administration course in March 2006 (see next item).

- In March 2006, the project team delivered a workshop entitled *QCDgrid System Administration Workshop 2006*. The event, which was hosted by the National e-Science Centre in Edinburgh, brought together system administrators and software developers from across the UK academic community and beyond, with the intention of providing training on the deployment, maintenance and development of QCDgrid-powered services.

  The workshop was funded by the GridPP programme, which provided financial support to cover the catering and administrative costs. Staff effort to prepare and deliver material for the workshop was contributed by the QCD Support project. Travel for a number of participants was covered by the UKQCD Travel Fund (PPARC Grant Number PPA/G/S/2002/00467). Computing resources and venue facilities were provided by the National e-Science Centre.

  Twelve people participated in the workshop, including system administrators involved with the UKQCD Grid, technical staff from the QCDOC team, and external participants from the HPC-Europa NA3 activity. Based on the feedback received from these participants, we are confident that the workshop was very successful, and the project team would hope to be able

to conduct a future workshop activity, following a similar structure to this one while building on the groundwork that has been done. In a future activity, the organisers would hope to be able to include presentations for QCDgrid users outwith the UKQCD collaboration.

06Q2 Comments

Project progress

During this quarter, the project team have continued to focus on Work Package 3 – developing the UKQCD contribution to the ILDG Metadata Catalogue (MDC). The web service implementation, previously developed in WP3.5, has been deployed onto the production UKQCD Grid web server during 06Q2. At the time of writing, metadata for eleven UKQCD production ensembles, including new data generated on QCDOC, is available through the service (Figure 3).

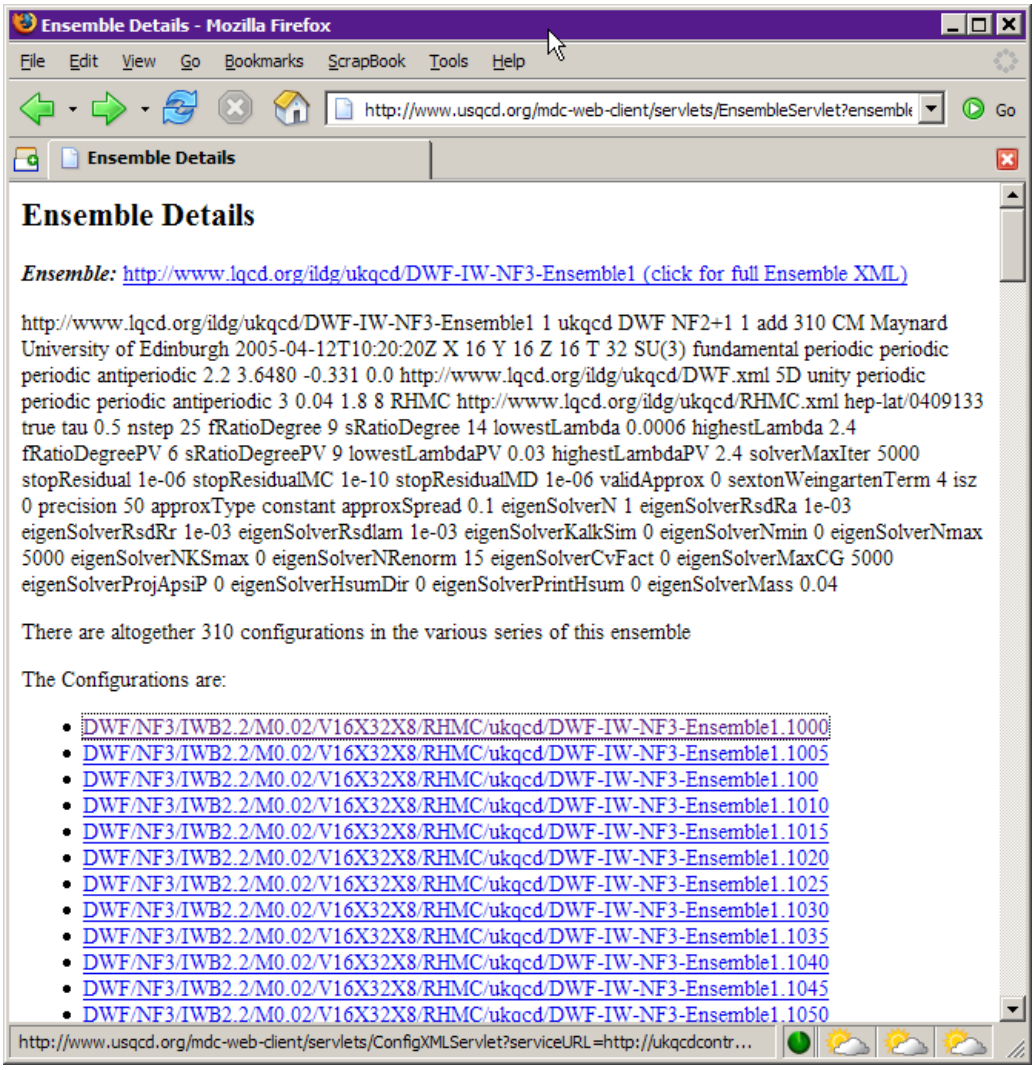

Figure 3: Screenshot of simple ILDG MDC web client displaying ensemble details for UKQCD dataset.

Since the production of the ILDG MDC was the first example of distributed software development within the ILDG working group, it has been pertinent to collate together the results and conclusions of the activity in order that any problems with the process can be identified and resolved, in preparation for the second phase of work on the ILDG File Catalogue. It was also an opportunity to gather feedback from early users of the service, in order to focus our remaining effort within Work Package 3. The output of the survey has been documented in D3.6. One of the key conclusions is that the most important need of the ILDG, at this time, is for a fully functional, easy-to-use client application to the MDC. To this end, the project team have replaced the original WP3.7, WP3.8 and

WP3.9 tasks with two new tasks (WP 3.7 and WP 3.8) that furnish a design and implementation of an ILDG MDC Client – based on the existing QCDgrid browser. Deliverable D3.7, the client design, was released in June. The implementation of this design is in progress and is expected in mid-July.

In May, the project team were pleased to announce the release of QCDgrid Version 1.4. This latest version includes several updates designed to simplify our contribution to ILDG, including support for:

- QCDML[4] Version 1.2.
- Dynamic configuration of client tools to use different XML schema for metadata.
- Separate hosting of the Metadata Catalogue and File Catalogue service components.

Version 1.4 also includes significant optimisations to the data replication process, plus a number of other bug fixes and improvements.

UKQCD Grid status

Following on from D1.5.2 UKQCD Grid Support Review 2005, a new, dedicated control node has been deployed onto the grid. In line with the recommendation of the review, the system is hosted within EPCC and is easily accessible to the QCDgrid team. All UKQCD Grid control services (Metadata Catalogue, File Catalogue, Data Replication Service, and Virtual Organisation Service) have been successfully migrated to the new host. A backup node will be configured at the University of Liverpool in the coming quarter.

A new storage element has also been added to the grid, located at the University of Edinburgh Advance Computing Facility. The new storage element has an initial capacity of 12 Tbytes and is linked directly to the QCDOC system, greatly simplifying the process of uploading new ensembles onto the grid. At the end of June 2006, the total capacity of the UKQCD Grid was approximately 80 Tbytes, hosting 40,000 dataset that occupy around 70% of the available space.

Dissemination activities

The project team are pleased to report two successful dissemination activities within this quarter:

- Firstly, our collaboration with the HPC Europa NA3 project, which began in 06Q1, has matured. The HPC Europa team have successfully deployed a QCDgrid instance for their project, with storage elements in the UK, Ireland and Denmark. This service is currently being tested by a selection of HPC Europa users. The QCDgrid team have established a regular bi-monthly meeting with the HPC Europa team and are liaising with them, regarding their usage and requirements for the system.
- Secondly, the project team have made initial contact with the Wellcome Trust Center for Cell Biology within the University of Edinburgh. The Cell Biology team are looking for a simple way to share large amounts of image data between multiple research centres across the UK. It is hoped that we can organise a meeting during July to discuss in detail their needs and the suitability of QCDGrid for their needs.

[4] QCDML is the ILDG specification for Lattice QCD metadata describing ensembles of gauge configurations. The specification is expressed as an XML application.

06Q3 Comments

Project progress

The most significant event during 06Q3 was the departure of Daragh Byrne from EPCC and hence from the QCDgrid project, on 25th August. Daragh has been replaced by Radoslaw Ostrowski, who joined both the project team and EPCC on 11th September.

Despite this change to technical staffing, the project team made good progress with the work plan during the period. The first version of the ILDG Browser (formerly referred to as the ILDG MDC Client) was released during August 2006 (see Figure 4 for an example of the application in action).

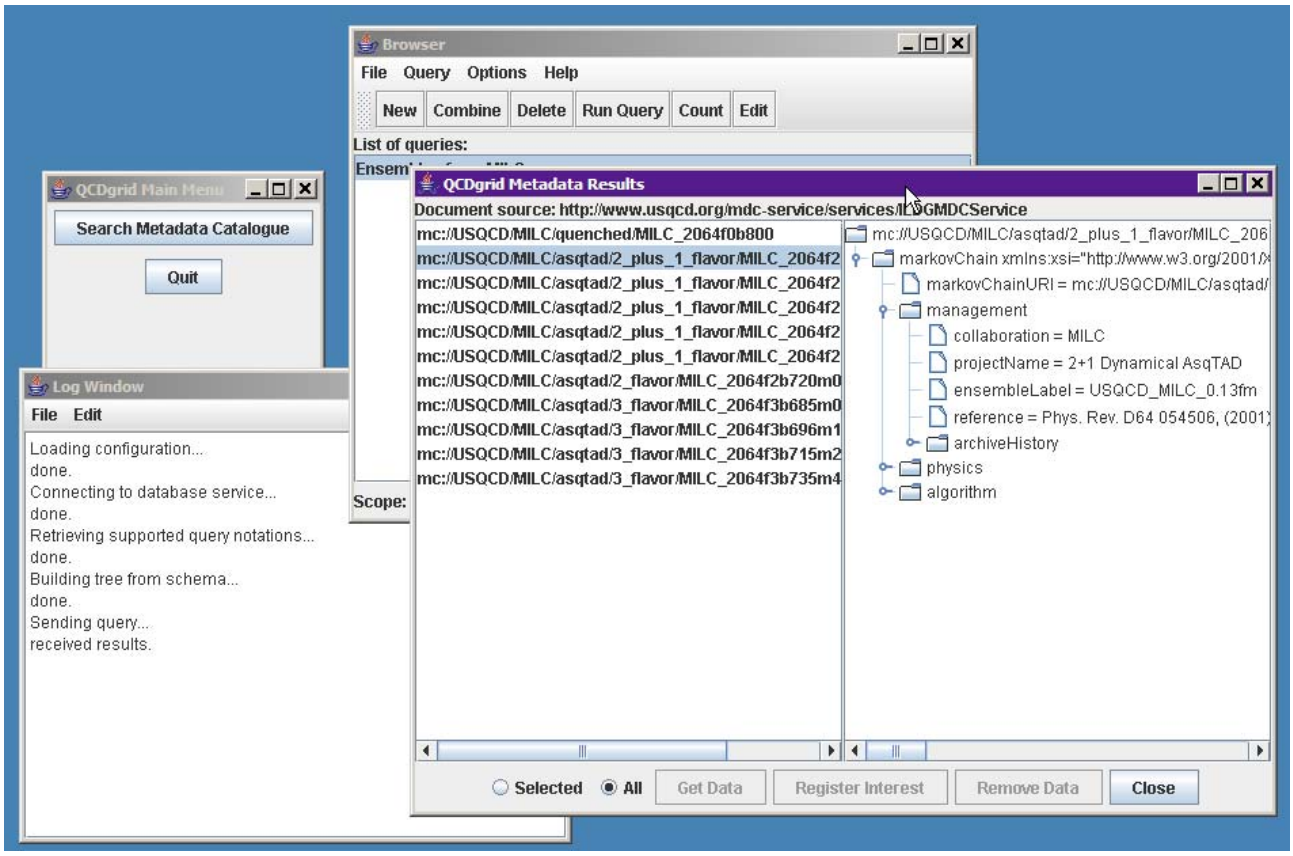

Figure 4: Example of ILDG Browser retrieving metadata pertaining to MILC configurations.

The ILDG Browser is available as a separate download, though is included in the QCDgrid software bundle.

Coinciding with the completion of the ILDG Browser implementation, a new version of the QCDgrid software (Version 1.5) was released in September. Alongside the ILDG Browser, this included a much-improved `qcdgrid-list` command, and a number of bug fixes.

Work has also begun on Work Package 4, which tackles the development of the ILDG File Catalogue. In July 2006, Daragh Byrne and George Beckett visited Dirk Pleiter, Michael Ernst and David Melkumyan in DESY Zeuthen (Germany). During this visit, the team made significant progress with Work Package 4 activities, including:

- Establishment of an ILDG Virtual Organisation built on top of the EGEE VOMS software (see https://grid-voms.desy.de:8443/voms/ildg/).
- Preparation of a draft ILDG VO Policy Statement, which has been submitted for approval to the ILDG Board.
- Initiation of a migration plan from the UKQCD certificate repository to the ILDG Virtual Organisation (see WP 4.1).

The team also performed a number of interoperability tests to determine the practicalities of transferring data between different collaboration's Storage Elements and learned about the current status of key EGEE SRM implementations. Also during the visit, the team progressed plans for a demonstration of the ILDG Grid at the Lattice 2006 conference in Tucson, Arizona.

UKQCD Grid Status

During the period, the UKQCD Grid has enjoyed good reliability and approximately 3,000 new configuration datasets (produced by the QCDOC super computer) have been added. These new datasets comply with the ILDG File Format specification, ensuring that they are immediately usable by other collaborations within the consortium.

Hardware maintenance has been required at both the Swansea site (which has encountered significant problems with RAID storage provision) and the Edinburgh site (where the Storage Element has been upgraded). Despite these incidents, the UKQCD Grid has continued to function effectively.

Dissemination activities

The project team are pleased to report that the ILDG presentation to the Lattice 2006 conference in Tucson, Arizona, was awarded the prize for *Best Scientific Poster*. The ILDG contribution to the conference included a presentation, poster, and demonstration of the ILDG Metadata Catalogue in action.

Secondly, the project team in conjunction with colleagues from the Wellcome Trust Center for Cell Biology, have secured funding from the eDIKT (e-Science Data Information and Knowledge Transformation) programme to conduct a feasibility study into the usage of QCDgrid software within the field of cell biology for the storage and distribution of microscopic image data. The feasibility study will be conducted during October—December 2006 and the final report should be available early in the new year.

Thirdly, Zhauyang Dong, who is a student of the MSc in High Performance Computing at the University of Edinburgh, has conducted an evaluation of the QCDgrid GUI Client/ILDG Browser for her dissertation[5] project. Her evaluation has focused on the general applicability of the browser functionality to non-Physics applications. Alongside the written report, the dissertation has also contributed additional functionality to the QCDgrid GUI Browser – specifically, additional Results Handler classes for tables and image data, plus several bug fixes. These additions are to be evaluated before inclusion into the next release of QCDgrid.

Finally, the OMII team have made an initial approach to EPCC to discuss the possibility of including the QCDgrid software into the OMII middleware family. The project team have little further information regarding the proposal at this time.

A new name for the software

To reflect the more general applicability of the QCDgrid software to scientific fields outwith Lattice QCD – as illustrated by the HPC-Europa project, the dissertation of Zhauyang Dong, and potentially by future work with Cell Biologists – the team have decided to rename the QCDgrid software. Following a competition held within UKQCD and EPCC, DiGS (Distributed Grid Storage) has been selected as the new name. This name will be applied to the next release of the software.

---

[5] Zhaoyang Dong, *Applications of QCDgrid*, MSc Dissertation, University of Edinburgh (2006). The report is expected to be publicly available before end of October 2006.

06Q4 Comments

Project progress

The QCDgrid project team have made good progress during the quarter. The team have completed work packages WP 4.1 and WP 4.2, which means the project has achieved Project Milestone 4.4.7. Most importantly, the work packages have advanced the development of two key components of the ILDG infrastructure:

1. An ILDG Virtual Organisation, aggregating user information and organisation-wide logistics into a centralised VOMS server
2. The ILDG File Catalogue that is used to track the location of Lattice Gauge Configuration datasets within the storage resources provided by regional grids such as UKQCD Grid.

The completion of these work packages was supported by the participation of George Beckett and Radoslaw Ostrowski in an ILDG Middleware workshop at Jefferson Laboratories, Virginia USA (10th—14th December 2006). The minutes from this meeting are available as the final deliverable to WP 3.1, bringing that work package to a conclusion.

UKQCD Grid Status

During the quarter, the number of files stored on the UKQCD Grid rose above 50,000 for the first time. The data includes new configurations produced by the QCDOC supercomputer, which are supported by QCDML metadata.

Reliability of the UKQCD Grid has been generally good during the period. Response time to failures has been improved thanks to new monitoring tools that have been adopted. Specifically, the team has deployed a modified version of the UK Grid Integration Test Script (GITS). A status page covering availability of basic grid services (GridFTP, GRAM, and RLS) and the amount of free space on UKQCD storage resources is maintained at:

http://ukqcdcontrol.epcc.ed.ac.uk:8080/gits/gits-results.html,

A sample screenshot is provided in Figure 5.

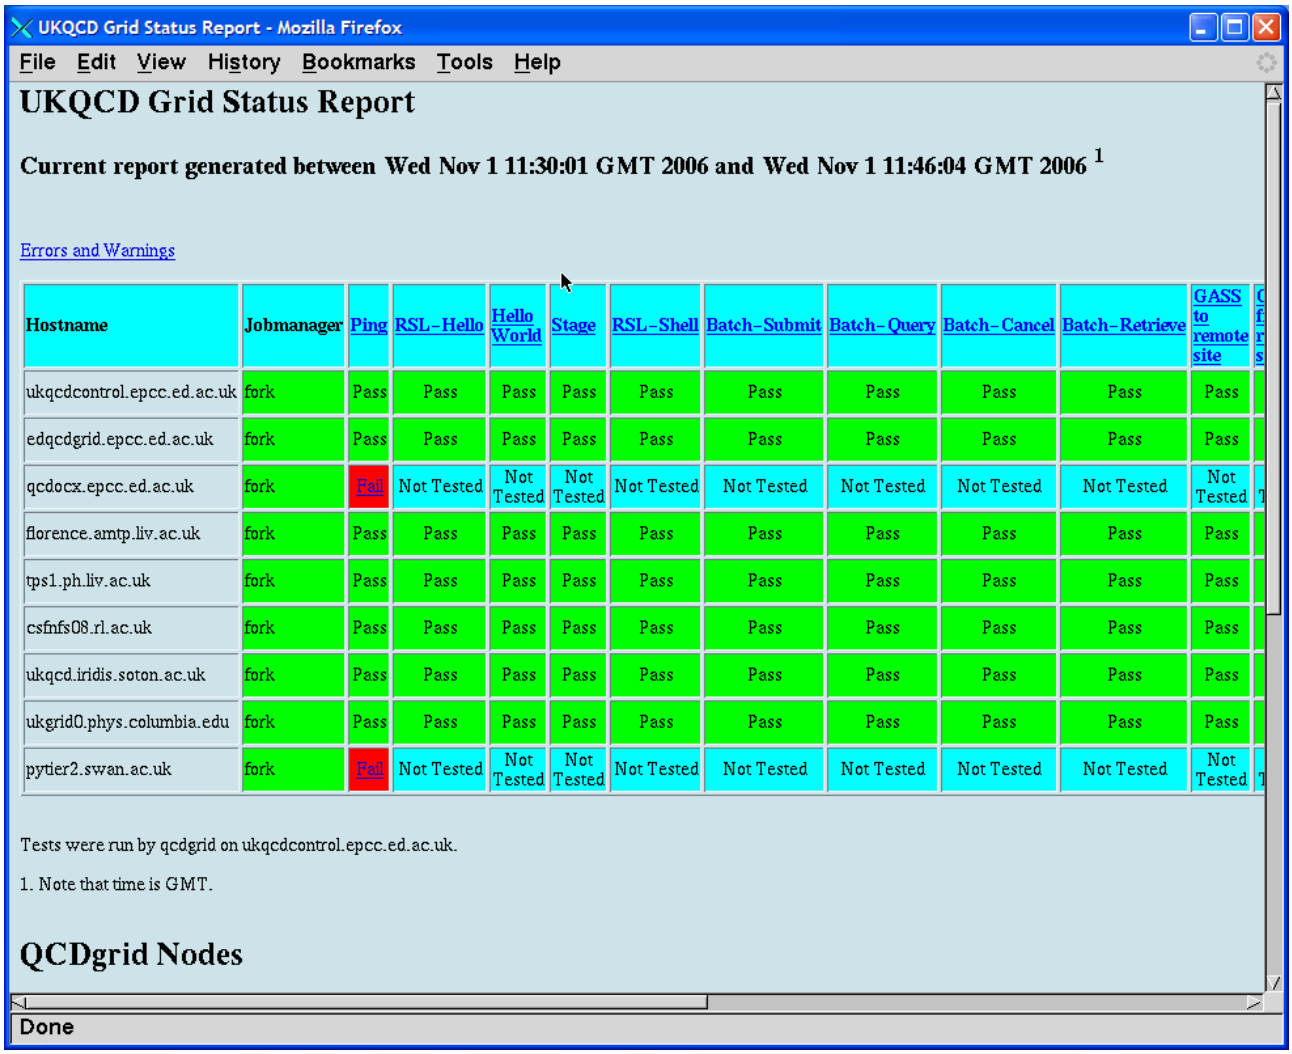

Figure 5: Sample screenshot of UKQCD Grid deployment of GITS monitoring tools.

Part of the work on the monitoring tools was contributed by an EPCC Summer School student, named Chris Daley.

Dissemination activities

During the period, the QCDgrid team participated in several dissemination activities, as follows:

- QCDgrid team participated in the workshop "VOMS for GridPP and the NGS", organised by the University of Manchester over Access Grid. The workshop featured presentations from Dr Mike Jones, Dr Alessandra Fort, and Dr Sergey Dolgobrodov.
- The project team contributed the UKQCD regional grid status report to the 9th ILDG Workshop (1 December 2006).
- On Thursday 7th December 2006, the project contributed a poster to the EPSRC International Review of ICT Research in the UK during the reviewer's visit to the University of Edinburgh. The poster highlighted project progress on bringing grid technologies and solutions to different scientific communities. A copy of the project is available on the project web pages (PDF).

The eDIKT-funded, one month feasibility study into the usage of QCDgrid software within the field of cell biology is close to completion. The final report for the activity will be available by the end of January 2007.

07Q1 Comments

Project progress

Project progress was good during the quarter. WP 4.3 has been completed, providing a security framework on which to build the ILDG interface to the data grid, as well as providing a valuable review of the existing UKQCD Grid security. The conclusions of WP 4.3 have been fed into WP 4.4, which is in progress and expected to be complete before the end of April 2007.

UKQCD have made an application to the PPARC Special Programme Grant to fund continued development of the UKQCD Grid beyond this project. The proposal focuses on the expansion of DiGS support for storage elements to include more recent standards such as Storage Resource Management (SRM). At the time of writing, UKQCD are awaiting the results of their application.

Status of UKQCD Grid

In January 2007, as part of our on-going efforts to maintain a reliable and performing data grid, we have upgraded the eXist database (used to host UKQCD metadata) from Version 0.9 to Version 1.1. Version 1.1 is a production release of the database, implying improved reliability and predictability. Furthermore, Version 1.1 has improved query performance, which will be of increasing benefit as the volume of metadata grows.

Several of the public interfaces to eXist changed in the move from Version 0.9 to Version 1.1. These, in turn, necessitated several modifications to the ILDG Browser client. As a result, Version 1.6 of the ILDG Browser client was released in March. This release also included a number of performance enhancements and bug fixes.

Dissemination activities

The final report of the eDIKT-funded feasibility study, which investigated the possibility of the DiGS software being utilised for data management in the field of cell biology, was released in February 2007. The main conclusion of the study is that DiGS is likely to have a role to play, when combined with an existing biology toolkit called the Open Microscopy Environment (OME) to provide a data distribution and replication functionality. Following on from the report, the Wellcome Trust Centre (in the University of Edinburgh) are deploying a test installation of OME, in order to gain more experience of the technology. George Beckett will then meet with the Wellcome Trust team in the summer (2007) to consider opportunities for further collaboration.

The project team are also undertaking a number of seminars to promote awareness of the work that the QCDgrid team have undertaken with the Lattice QCD community, as well as to help us identify new applications for DiGS software. The first of these seminars was given at the National e-Science Centre in Edinburgh during March 2007. Subsequent talks are planned for the eDIKT Technical Workshop in June 2007 and the EPCC Software Development Group in April 2007.

07Q2 Comments

Project progress

Project progress on the design and implementation of the ILDG File Catalogue has been slower than expected during the period. The design document (D4.4) was released in early May – which was five weeks later than planned. This has contributed to a subsequent delay in the completion of WP 4.5, which remains on-going. The main source of the delay (for both deliverables) has been problems with our chosen proxy delegation service – the gLite Delegation Service (D/S):

- A number of problems were encountered during the design phase, which included an investigation of gLite D/S. However, with the help of the gLite D/S development team (which is led by Dr Andrew McNab, University of Manchester), we were able to close our investigation and thus complete WP 4.4.

- During the implementation phase (WP 4.5), a serious problem has been identified in the current version of gLite D/S. Specifically, both ourselves and our colleagues in DESY Zeuthen have been unable to get proxy certificates – generated by the service – to be accepted by either other gLite services (as used by the LaTFoR grid in Germany) or Globus services (as used by UKQCD Grid). Discussions are on-going with the gLite D/S development team, though the problem has not been resolved.

  We have therefore begun contingency work, using an earlier release of gLite D/S (Version 3.0), which has been confirmed to generate valid proxy certificates. Unfortunately, this earlier release of gLite D/S is significantly less functional and has a different interface. Additional implementation time (around 2—3 weeks) is required to address these discrepancies.

The delay has placed our plans, to demonstrate the ILDG File Catalogue at the Lattice 2007 conference (Project Milestone 4.4.9), at risk. We have taken three steps to attempt to control this risk:

- We have shared the implementation work on the File Catalogue web service between the QCDgrid team and our colleagues at DESY Zeuthen. To this end, Radek Ostrowski visited DESY Zeuthen between 4th—6th June to work on the task. This was a successful visit and a candidate implementation of the File Catalogue web service was completed during the visit.
- We have allocated effort to configure the QCDgrid project development system as an ILDG resource – as the development system does not host real scientific data, the time required to deploy ILDG services on the system is significantly less than for the production UKQCD Grid.
- If necessary, we propose to de-scope the functionality of the ILDG client tools developed in WP 4.6 and WP 4.7.

In parallel with this work and as part of on-going preparation for the Lattice 2007 conference, an abstract has been submitted (under WP 5.1):

> ***Towards an interoperable International Lattice Datagrid***
>
> *P. Coddington, G. Beckett, N. Ishii, B. Joo, D. Melkumyan, R. Ostrowski, D. Pleiter, M. Sato, J. Simone, C. Watson, S. Zhang for the ILDG Middleware Working Group.*
>
> *The International Lattice Datagrid (ILDG) is a federation of several regional grids. Since several of these grids have reached production level, an increasing number of lattice scientists start to benefit from this new research infrastructure. The Middleware Working Group has the task of specifying the ILDG middleware such that interoperability among the different grids is achieved.*
>
> *We will present the current status of our work. A live demo will allow participants to test access to the vast amount of gauge configurations which have become available through ILDG.*

We have received confirmation from the conference organisers that this abstract has been accepted.

Status of UKQCD Grid

Other project activities have progressed more smoothly.

- During May, the team completed the migration of user management in the UKQCD Grid to a VOMS-based system. This replaced an existing LDAP service (developed by the EDG VO project) that had been in use since 2003.

  At the time of writing, a single VOMS server is available, hosted in DESY Zeuthen. This is administered by a VOMS-RS service hosted in DESY Hamburg https://grid-voms.desy.de:8443/vo/ildg/vomrs. In the future, it is hoped that additional (mirror) servers will become available. To this end, the QCDgrid project team are working with Alessandra Forti and

Sergey Dolgdrobov (GridPP and University of Manchester) to set up an ILDG VO on the GridPP test VOMS server.

- During June, Craig McNeile (a UKQCD member from University of Glasgow) set up a new client node to the UKQCD Grid (on ScotGrid). This new client node is being tested, at the time of writing, though once in production it should provide a single point of access for computational resources on ScotGrid and storage resources on UKQCD Grid.

Dissemination activities

Building on the work of the previously reported, eDIKT-funded feasibility study, which investigated the possibility of the DiGS software being utilised for data management in the field of cell biology, George Beckett and Russell Hamilton (Dept. of Biochemistry, University of Oxford) have submitted a proposal to the STFC Bio Mini-PIPSS call (http://www.stfc.ac.uk/KE/FOpp/PIPSS/Calls/BioMini.aspx). The purpose of this project is to enable distributed scientific collaborations in the field of Cell Biology to share biological image data in a simple, secure, and effective manner, leveraging the experience and technology developed within the QCDgrid project. We are expecting the receive a decision on our proposal within 2007Q3.

George Beckett also continues to work, with the support of the eDIKT programme, on disseminating the work of the QCDgrid project. In June 2007, he led a technical workshop on "Data Management" using UKQCD Grid as a key illustration of the topic.

07Q3 Comments

Project progress

Progress during this quarter has been good, with two significant milestones to report:

- The project team (as members of the ILDG) successfully demonstrated a functioning International Lattice Data Grid, at the XXV International Symposium on Lattice Field Theory (Lattice 2007), in Regensburg, Germany. This demonstration included the ILDG Metadata Catalogue (see WP 3) and the ILDG File Catalogue (see WP 4), plus Storage Elements (also covered in WP 4). The demonstration integrated regional grids from five collaborations – that is: ourselves in UKQCD; CSSM (Australia); JLDG (Japan); LatFor Data Grid (Germany); and USQCD (USA). The demonstration is supported by a paper that has been submitted to the conference proceedings. This fulfils project milestone PM 4.4.9.

- The project team have released DiGS Version 2.0. This is a major update to the QCDgrid software, which includes numerous new and extended functionalities. The primary objective in developing DiGS Version 2.0 has been to allow UKQCD data resources to be integrated into ILDG. However, the team have also taken the opportunity to implement changes in order to satisfy more significant user requirements, plus to generalize the functionality of the technology, improving its applicability to other scientific applications. Full details of the new release are provided in the news item on NeSCForge.

  Because of the two month delay that was reported in 2007Q2, we were unable to deploy DiGS software onto the UKQCD Grid in time for the Lattice 2007 conference. To mitigate this, we contributed to the Lattice 2007 demonstrator using candidate release (DiGS 2.0 RC1), which was deployed onto the UKQCD Development Grid in Edinburgh.

  The release of DiGS Version 2.0 contributes to the tracker milestone PM 4.4.3.

Looking to the future, the project team have moved their attention on to Work Package 6, beginning work on WP 6.1 in early September. With regard to the work plan, we remain two months behind schedule, and will establish the impact of this on Work Package 6 in the 2007Q4 period.

Status of UKQCD Grid

The QCDgrid team have been liaising with the GridPP systems team (primarily, Matt Hodges at RAL), regarding the migration of data currently held on a GridPP storage element at RAL to a new host. This new host, which is also at RAL, will provide increased capacity (16 Terabytes) to the UKQCD Grid, and is expected to be the last WLCG "classic" storage element incorporated into the grid: the tasks defined in Work Package 6 are intended to allow the next release of the DiGS software to interact with SRM-based storage elements.

The QCDgrid team (as members of ILDG) have continued to liaise with Alessandra Forti and her team, in Manchester, to publish ILDG Virtual Organisation (VO) information onto the GridPP VOMS service – as reported in the 2007Q2 report. The work was suspended by DESY Hamburg in early July, pending the arrival of several important security fixes for the VOMS-RS code base. These were implemented in mid-September, so work will continue on this task in 2007Q4.

Following on from the release of DiGS Version 2.0 at the end of September, we have begun rollout of the new software onto the UKQCD Grid. This rollout is expected to be completed by mid-October.

Dissemination activities

The most notable news is that the proposal to the STFC Bio Mini-PIPSS call, which was noted in the previous quarter, has been accepted, with one year of funding (1 FTE) to enable distributed scientific collaborations in the field of Cell Biology to share biological image data in a simple, secure, and effective manner, leveraging the experience and technology developed within the QCDgrid project. This funding offers the opportunity to continue the development of software and expertise in distributed data management, and will ensure future support and development of the QCDgrid/DiGS software.

During 17th—20th September, Dirk Pleiter (DESY Zeuthen) visited the School of Physics in Edinburgh. This provided an opportunity for the QCDgrid team to meet with Dirk to discuss on-going ILDG activities, and consider opportunities for future funding. Building on this visit, George Beckett and Dirk Pleiter (as representatives for the ILDG Middleware Group) submitted a set of usage cases for ILDG, to the European Grid Initiative (EGI) call issued by Jan Kmunicek, Masaryk University, Czech Republic.

5. Meetings & Papers

5.1 List of Conference Papers

- M.G. Beckett, P. Coddington, N. Ishii, B. Joó, D. Melkumyan, R. Ostrowski, D. Pleiter , M. Sato, J. Simone, C. Watson, S. Zhang, *Towards an interoperable International Lattice Datagrid*, Submitted to Proceedings of XXV International Symposium on Lattice Field Theory (Lattice 2007), in Regensburg, Germany (July 2007).
- C.M. Maynard (University of Edinburgh), D. Pleiter (NIC, Zeuthen & DESY, Zeuthen), *QCDML: First Milestone for Building an International Lattice Data Grid*, EDINBURGH-2004-18, Sep 2004. 9 pp. Presented at 22nd International Symposium on Lattice Field Theory (Lattice 2004), Batavia, Illinois, 21-26 Jun 2004. Published in Nucl. Phys. Proc. Suppl. 140:213-221, 2005.

5.2 List of Conference Talks

- 13th June 2007: eDIKT Technical Workshop. Discussion session on scientific data management, led by George Beckett.
- 7th March 2007: Seminar at National e-Science Centre. Presentation by George Beckett entitled "DiGS: A data grid for UKQCD and beyond".
- 23rd—28th July 2006: The XXIV International Symposium on Lattice Field Theory in Tucson, Arizona. Poster presentation by Chris Maynard (University of Edinburgh and ILDG) and Balint Joó (Jefferson Laboratories and ILDG), entitled "Progress in building the International Lattice Data Grid".
- 28th March 2006: HackLatt06, National e-Science Centre in Edinburgh, UK. Presentation by Daragh Byrne entitled "The QCDgrid Software Suite".
- 9th—13th May 2005: *PPARC UK e-Science Postgraduate School*, Edinburgh, UK. Presentation by Chris Maynard entitled "QCDgrid, a Grid for UKQCD". The audience consisted of PPARC-funded, postgraduate students (typically in their second year of study). An electronic copy of the presentation can be obtained from the above link.
- 13th—14th April 2005: *UKQCD HackLatt 2005 workshop* in Edinburgh, UK. Two presentations from Daragh Byrne and James Perry, entitled "The QCDgrid software suite, an introduction" and "QCDgrid administration", respectively.
- 21st—24th September 2004: *Lattice QCD simulations via International Research Network* in Shuzenji, Japan. Presentation by Chris Maynard entitled "The current working status of QCDGrid in UK".
- 16th September 2004: PPARC Workshop *Technology Opportunities from CERN: The Impact of Big Physics on Industry* (see http://www.pparc.ac.uk/in/cern50.asp for more information). Poster presentations plus demonstration by James Perry and George Beckett.
- 14th—15th September 2004: *11th GridPP Collaboration Meeting* at the University of Liverpool (see https://www.gridpp.ac.uk/gridpp11/ for more information). Two presentations were provided by Craig McNeile (University of Liverpool) and James Perry, respectively.

5.3 List of publications

- The GridPP Collaboration (197 authors), *GridPP: development of the UK computing Grid for particle physics*, J. Phys. G: Nucl. Part. Phys. 32 (2006) N1-N20.

- C.R. Allton, W. Armour, M.G. Beckett, D. Byrne, C.T.H. Davies, S. Downing, J.M. Flynn, A.C. Irving, A.N. Jackson, B. Joó, R.D. Kenway, C.M. Maynard, C. McNeile, J. Perry, L. Smith, Z. Sroczynski and A.S. Trew, *QCDgrid: A Grid Resource for Quantum Chromodynamics*. Journal of Grid Computing 3, 2005, pp. 113 – 130.

5.4 Other Dissemination Activities

- M.G. Beckett and D. Pleiter *International Lattice Data Grid, Use Cases for Deployment and Usage*, Contribution to European Grid Initiative (September 2007).
- M.G. Beckett and R. Hamilton, *DiGS, A Data Grid for Cell Biology*, eDIKT Project Report (February 2007).
- 20th—21st October 2004: *Visit by Olivier Brand-Foissac* (of the Université de Paris Sud in Orsay, France), acting as a representative for the new French component of the ILDG collaboration.

5.5 Other Project Publicity

- 7th December 2006: *DiGS – Distributed Grid Storage for UKQCD*, poster contributed to the EPSRC International Review of ICT Research in the UK.
- Horst Simon (Chair of International Review Panel), *International Review of Research Using HPC in the UK*, EPSRC and Deutsche Forschungsgemeinschaft (December 2005). The review included very positive comments about the UKQCD consortium and the project. See 05Q4 Comments in Section 4 for more information.
- 3rd June 2005: *GridPP news article* on UKQCD HackLatt 2005 workshop.

6. Effort Delivered

Effort figures are provided within the accompanying spreadsheet, named "UKQCD_LogBook_Sep_2007_effort_figures.pdf".