

Progress in building the International Lattice Data Grid

B. Joó*

Jefferson Lab, Newport News, VA 23606, USA

E-mail: bjoo@jlab.org

C.M. Maynard

EPCC, School of Physics, Univeristy of Edinburgh, Edinburgh EH9 3JZ, UK

E-mail: c.maynard@ed.ac.uk

ILDG Metadata and Middleware Working Groups

The International Lattice Data Grid (ILDG) [1] has been making steady progress from an abstract concept to a more concrete implementation. The last year saw the completion of the MDC Web service specification, and the ILDG member groups; CSSM (Australia), Latfor data grid (LDG) (Germany), LDA (Japan), QCDgrid (UK), USQCD (USA), have either completed or are moving towards its implementation The Replica Catalogue service is in the process of being defined and data transfer trials have begun.

XXIV International Symposium on Lattice Field Theory

July 23-28 2006

Tucson Arizona, US

*Speaker.

1. Introduction

Many collaborations are storing and accessing data and other resources via grid technologies, having built grids on a national or collaboration basis. These regional grids are now well established. It is the goal of the ILDG to make these grids interoperable.

Lattice QCD computations are complex and expensive. Increasingly groups are sharing the cost of generating ensembles of configurations directly by forming collaborations of collaborations, or indirectly by developing open source software. A natural progression make data open source, that is, data sharing. How the lattice community shares data is both a sociological problem, *i.e.* what data is available and who can access it, how is credit assigned to individuals, and a technological one, *i.e.* what is and how to access this shared data.

The key to solving the technological problem is the use of standards. The data has a defined standard format, and the description of that data; the metadata has a defined and extensible scheme. The services that access and interact with this data have a defined *interface*, but the implementation of these services is left to the individual groups. In this way individual groups retain control over their resources, but remote users can access remote data using their local systems.

Two services have so far been defined: the Metadata Catalogue (MDC) and the Replica Catalogue (RC) [1]. These specify the interfaces between the different grids. A schematic diagram of the service architecture is shown in figure 1. The MDC service has now been implemented by several groups. This is demonstrated using two prototype clients for reading this service. Member groups have begun the work of implementing the RC. Data transfer trials are also underway.

The starting point for the ILDG is sharing gauge configurations. The common format for this data is based around the SciDAC LIME formatted files [2]. LIME is a limited version of the Direct Internet Message Encapsulation (DIME) which is now part of the Microsoft .NET framework. LIME files are comprised of a series of messages, each message being made up of one or more records. The ILDG format specifies records to describe the data layout, to hold the binary data, and to hold the Logical File Name (LFN) of the gauge configuration. Other internal or private records can be freely added to the file by producing collaborations to annotate their data as desired. Such extra records may be ignored by ILDG tools.

Metadata is data about data. In order to successfully share data it is critical to know what the data is. The obvious way of defining an ensemble of gauge configurations is mathematically, by writing down the action. However, an ensemble of gauge configurations has many different aspects besides the action, such as the size of the system, and the algorithm, software and hardware that were used to generate it. Much of this information is hierarchical. It is natural to use a mark-up language for the metadata, the obvious one being XML. An XML application (QCDml) has been defined. This is a set of rules for individual XML documents, so that a computer application can process the metadata. This might be a search query, extracting a certain element, or insertion or retrieval from a database.

Defining this XML schema was a balance between completeness (including as much as possible), extensibility so that future extensions to the schema do not require existing XML documents to be re-written and useability, so the burden of creating these documents does not prevent groups doing so.

1. ILDG data format for gauge configurations has been defined.
2. QCDm1.3 is the most current version of schema for describing metadata.
3. Gauge configurations in ILDG format, with QCDm1.3 metadata now being uploaded to regional grids.
4. ILDG MDC v1 service, behaviour and standard compliance tests have been defined.
5. The interoperation between local clients and MDC services has been successfully verified.
6. An ILDG V/O has been created.
7. The ability of users to down-load files using SRM and Gridftp from both the LDG and the storage node at FNAL has been verified. Access to data at these sites requires membership of the ILDG V/O.

2. Architecture

The architecture of the ILDG is sometimes described as a “grid-of-grids”. While this is a useful concept, it is somewhat limited. In reality, many different services from remote sites allow access to these resources through common interfaces. This is shown in figure 1. Below we define what is meant by the terms in the diagram.

LFN: Logical File Name. This the unique name for a file, belonging to a *namespace*. Each collaboration controls it’s own namespace. The ILDG namespace is an aggregation of the these namespaces.

RC: The Replica Catalogue (also called File Catalogue) maps the LFN to individual file instances. Each collaboration has it’s own RC, and these are combined using a common interface. It is intended that the ILDG RC service will be able to (though not required to) track replicas of on remote systems.

MDC: Metadata Catalogue: Each regional grid provides one MDC service for querying its metadata. This service is accessed by a web service client (web page, command line tool, or a graphical user interface (GUI) application). Due to the uniform and standard interface of the MDC Service clients can easily query individual or multiple web services (aggregation).

FTS: File Transfer services: The Middleware Working group has adopted the Storage Resource Manager (SRM v2) to manage the transfer of files. This service can stage files from tape and negotiate transfer protocols (http, gridftp, srmget, wget etc). However, as SRM v2 is not universally available at all sites, one can access files by connecting to the storage element file servers directly.

Security Certificate: Access is authenticated with X509 certificates which are issued by a Certificate Authority (CA). These can be used to identify both users and machines. Certificates allow “single sign on” to grid services. User access may be delegated to web service interfaces.

Authentication Service: A Virtual Organization (VO) is a group of people using common resources. Resource providers make their resource available to the VO. VO administrators decide who is a member. The Virtual Organisation Membership Service (VOMS) allows for common administration of the VO. The VOMS service will facilitate the implementation of access policies.

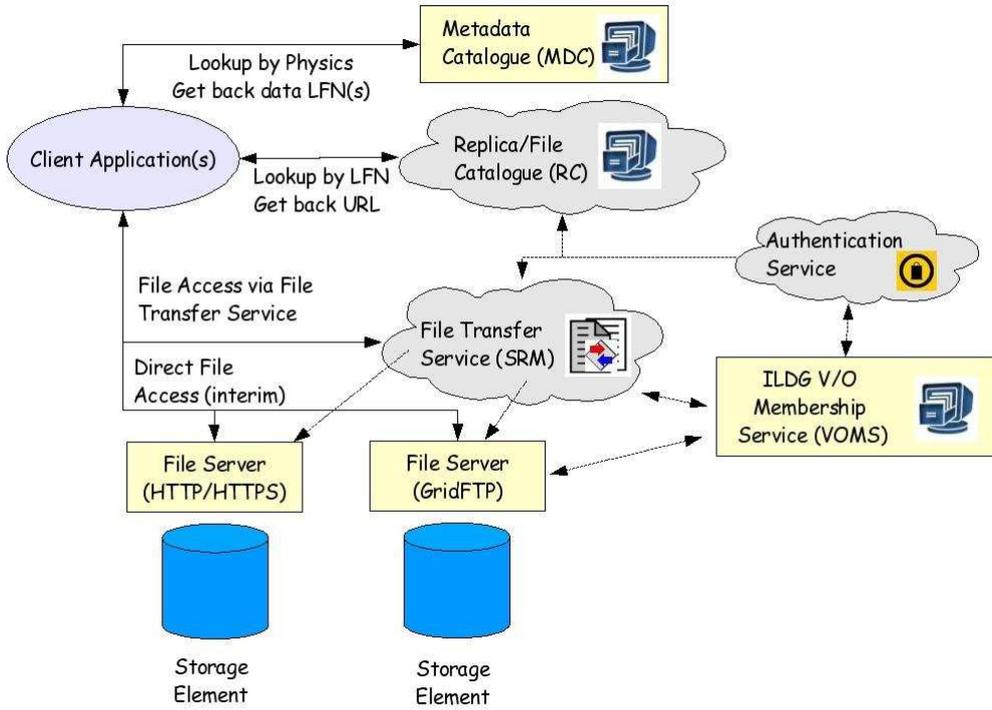


Figure 1: Figure 1. A schematic diagram of the architecture of the ILDG.

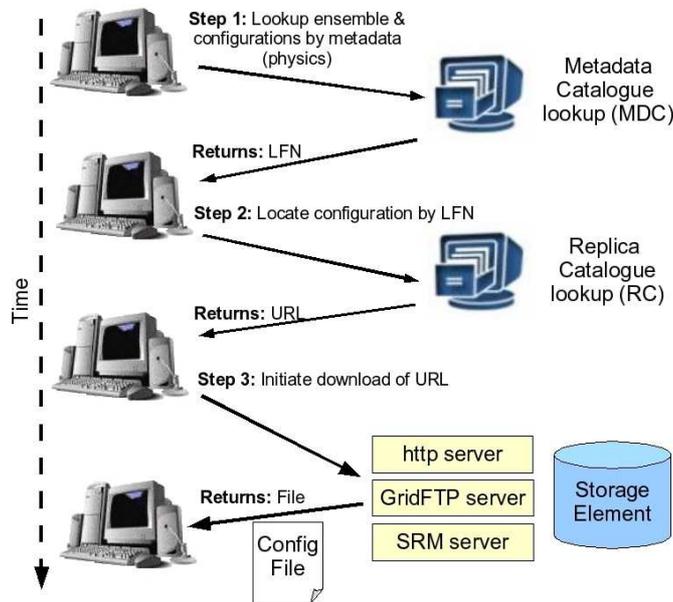


Figure 2: Schematic diagram of how access data that has been made available to the ILDG.

POS(LAT2006)028

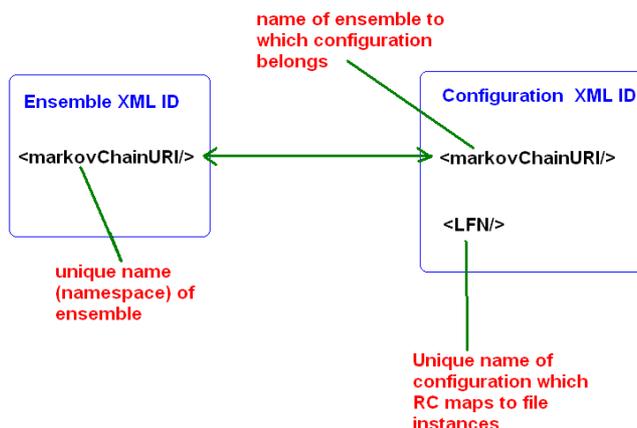


Figure 3: Schematic diagram showing how the ensemble and configuration XML (IDs) are related to each other.

3. Metadata

The metadata is divided into that which is generic across an ensemble, such as the action, and that which is specific to each configuration such as the value of an observable, such as the average plaquette, on that configuration. Correspondingly separate schemata have been defined for ensembles and configurations. This has several advantages. First, it reduces the burden on marking up configurations. The ensemble metadata might be created by some human intervention, and only needs to be done once per ensemble. The configuration XML contains metadata that the application creating the data already knows, and thus it is relatively easy to write to a file. Secondly, since metadata queries are typically searches for ensembles rather than individual configurations, this factoring of metadata reduces the number of frequently accessed instance documents which a database has to search through. This is important when one considers that hierarchical, rather than relational databases may not scale very well.

The configuration and ensemble metadata are linked together by a common XML element. This is shown in figure 3. The *Markov Chain URI* (MCU) in the `<markovChainURI/>` tag, is the unique name in the ILDG namespace for the ensemble. Each configuration belonging to this ensemble has the same name for its MCU. The configuration XML has the LFN, which can then be sent to a RC to get an actual file instance. For verification that the data is correct, the value of the average plaquette is part of the configuration XML, as is a CRC checksum.

Submitting queries to the MDCs can be done in a number of ways. For instance figure 4 shows the graphical user interface for building Xpath queries to search the QCDgrid MDC client. This tool has been adapted to submit queries to all ILDG MDCs. Our live demonstration is of two prototype MDC web clients from USQCD [4] and LDG [5]. The MDCs of all the ILDG sites can be searched using either client. The returned XML is rendered into a web-page.

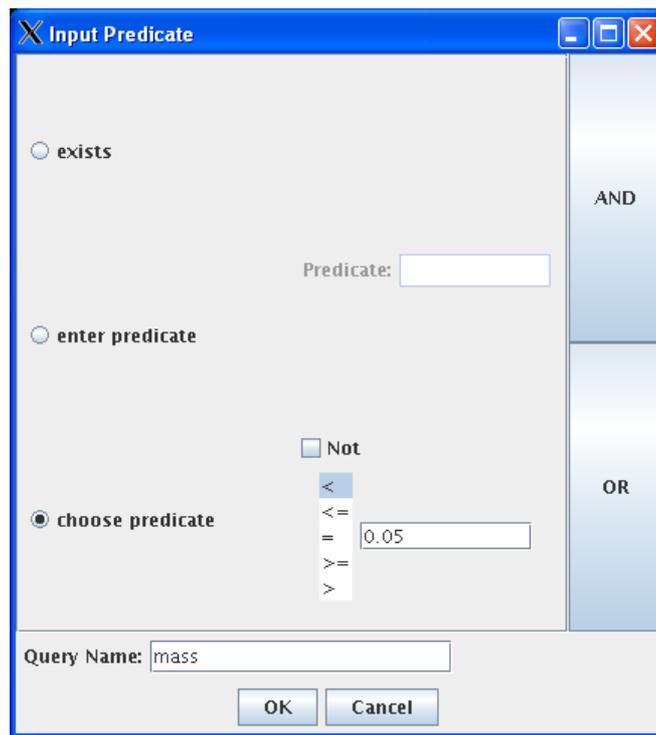


Figure 4: Screen shot from QCDgrid [3] GUI for graphically building Xpath query.

4. Acknowledgements

We thank the many people who have worked on making the ILDG a reality, including (but not limited to), the Middleware Working Group [1] and the Metadata Working Group [1]. The UKQCD grid project, QCDgrid, is a gridPP project [3]. LDG [5] is the Latfor project. U.S. funding for work on ILDG comes from the SciDAC programme of the U.S. Department of Energy.

References

- [1] <http://www.lqcd.org/ildg>
- [2] <http://www.lqcd.org/scidac>
- [3] <http://www.gridpp.ac.uk/>
- [4] <http://www.usqcd.org/mdc-web-client/index.jsp>
- [5] <http://www-zeuthen.desy.de/ape-cgi-bin/ildg-mdc.cgi>