International Lattice Data Grid, Use Cases for Deployment and Usage

M.G. Beckett, D. Pleiter

28 September 2007

In this paper we give a short overview on the International Lattice Data Grid (ILDG) and discuss a few use cases that provide hints on how ILDG would benefit by improving the European grid infrastructure.

1 Background

The ILDG is an international community of computational particle physicists who want to facilitate the sharing of data to accelerate progress in the field of Lattice Quantum Chromodynamics (LQCD).

The community was formed in 2001 and includes regional collaborations in: Australia (at CSSM), mainland Europe (Latfor Data Grid), Japan (Tsukuba), UK (UKQCD) and USA (BNL, Fermilab and JLab).

The datasets that the community wish to share are called Lattice Gauge Configurations (or simply 'configurations'). These are typically generated in families (called 'ensembles'), with an individual ensemble containing possibly hundreds of individual configurations.

Ensembles are expensive to generate (requiring HPC-type compute resources), and volumous to store (with ensembles typically consuming hundreds of Gigabytes). Well-established collaborations have already provisioned substantial HPC resources, and allocated Terabytes of storage capacity (both disk-based and tape-based) on which to host this data. These compute and storage resources are often administered by a third-party provider (such as GridPP in UK). Because of this, the regional grid is often restricted in the software/services that they can deploy onto the resources.

Since ILDG is not in a position to develop and deploy a common infrastructure across the different regional collaborations, an alternative approach has been pursued, and an ILDG Middleware Working Group has been set up to design and implement an interface

layer on top of, and uniting, the diverse infrastructures that are in place in different collaborations (or regional grids) to form what is called a 'grid of grids'. The underlying regional grid infrastructures are build on top of: grid middlewares such as the Globus Toolkit and gLite application stack; or bespoke web service components and applications.

From the ILDG layer, the 'grid of grids' consists of four components:

- A Metadata Catalogue, which holds scientific descriptions of published configurations and ensembles. In practice, the Metadata Catalogue is assembled from individual catalogues provided by regional grids and exposed through a web service layer.

- A File Catalogue, which maintains mappings between the unique and persistent identifier for each dataset and a list of one or more SURLs (internet locations) at which the dataset can be found. Again, in practice, the File Catalogue is distributed across regional grids, and exposed to an ILDG user through a web service layer.

- Storage Elements, which are internet-accessible resources on which datasets are hosted. Storage elements are provided by regional grids. GridFTP and HTTP are the two transport protocols that are supported. Furthermore, some regional grids organise storage elements behind an SRM-compliant service.

- A Virtual Organisation Manager, which hosts a list of ILDG members. Each member record includes basic information such as the member's regional grid affiliation, allowing regional grid to apply a course level of access control to local and non-local ILDG collaborators. Membership administration is co-ordinated through a single VOMS-RS that is, at the time of writing, hosted by Latfor Data Grid. VOMS-RS pushes membership information to one or more VOMS servers, from which regional grids can retrieve the information and thus populate their own authentication/authorisation services.

2 Sample Use cases

2.1 Use Case ILDG1

Title

Registering to use ILDG.

Description

The agreed process by which a LQCD researcher may register to access datasets held on ILDG.

Actors involved

- LQCD 'researcher' (from a participating regional grid) that wishes to join ILDG.
- An appointed 'representative' for particular regional grid to which researcher belongs.

Pre-conditions

Researcher has a valid X.509 certificate, issued by any of the Certificate Authorities of the International Grid Trust Federation (IGTF). This must be installed in the researcher's web browser client.

Steps

1. Researcher establishes authenticated connection to VOMS-RS and completes the registration form. The researcher is required to identify their home regional grid and nominate a representative to approve their application, from a list of predefined representatives for the regional grid.
2. Researcher completes an e-mail verification process, at which point they become a candidate for membership of ILDG.
3. Researcher requests to join a regional grid group, again through the VOMS-RS interface. Researcher needs to accept the terms of the VO Usage Rules at this step.
4. The nominated representative is notified of the application, and either accepts/rejects the application, having confirmed that the researcher is: (i) a genuine LQCD scientist, (ii) affiliated to their particular regional grid.
5. A description of the decision (in Step 4) is circulated to all representatives (from every regional grid). We assume application is accepted as if rejected the process stops at this step.

6. The researcher's information is automatically pushed to all VOMS systems that host ILDG VO information, by VOMS-RS.

7. During the next scheduled synchronisation operation (usually within 24 hours) regional grid authentication and authorisation services retrieve the new user information and update their local user information accordingly.

8. After no more than 24 hours, the researcher is able to access the resources of ILDG with privileges that correspond to their role in the community.

Middleware involved

- VOMS-RS, administrative interface for the application.
- VOMS, service (possibly more than one instance) which hosts current ILDG membership information.
- Regional grid authentication and authorisation services, such as GSI.

Applications involved

No non-standard applications are involved.

2.2 Use Case ILDG2

Title

Obtaining a valid X.509 certificate

Description

To become member of the VO ILDG a LQCD researcher has to have a valid X.509 certificate, issued by any of the Certificate Authorities of the International Grid Trust Federation (IGTF).

Actors involved

- LQCD 'researcher' (from a participating regional grid) that wishes to obtain a X.509 certificate.
- A Certificate Authority (CA) or corresponding Registration Authority (RA).

Steps

1. Researchers generates private key and certificate request and submits the latter to the CA.
2. The researcher is identified by the RA.
3. Certificate is sent by CA to researcher.

Notable items

At the time of writing, Registration Authorities are not available at all European sites where research on LQCD is carried out.

2.3 Use Case ILDG3

Title

EGEE compliant storage elements at HPC centres

Description

Members of the VO ILDG produce vast amounts of data on supercomputers installed at national (and in future: European) supercomputer centres.

Actors involved

- LQCD 'researcher' that uses supercomputing resources and wants to: (i) upload generated data to a local grid-enabled storage element; (ii) download data from a grid-enabled storage element for further processing.
- Supercomputing centres providing both compute and storage resources for LQCD applications.

Steps (i)

1. Generate data on supercomputer
2. After job completion upload data to local grid-enabled storage element

Steps (ii)

1. Pre-stage data on local grid-enabled storage element.
2. Once job execution has completed, download data to local disk and process data.

Middleware involved

- Storage element software, e.g. dCache
- Data management client tools, e.g. srmcp

Applications involved

Non-standard applications are not used during access of grid resource.

Notable items

At the time of writing, only a subset of regional grids have provide access to the local storage resources via EGEE-compliant interfaces (in particular, SRM).

2.4 Use Case ILDG4

Title

Data management client operations on different platforms

Description

The use of gLite software is essentially restricted to a small set of Linux flavours. For a data grid like ILDG it would be of major benefit to have data management client tools available for a larger set of unix platforms.

Steps

1. Perform sign-on, i.e. generate proxy certificate.
2. Manipulate content of file catalogue.
3. Up- or download data from a storage element, delete files from a storage element, replicate data between storage elements.

Middleware involved

- Tools for managing proxy certificates
- Data management client tools (e.g. gridftp, srmcp, GFAL, lcg-util)

Applications involved

No non-standard appications are involved.

2.5 Use Case ILDG5

Title

Search for scientific data

Description

An LQCD researcher is looking for an ensemble of gauge configurations that exhibit specific scientific properties.

Actors involved

An LQCD 'researcher'.

Prerequisites

Researcher is registered member of ILDG VO.

Steps

1. Researcher opens the ILDG Browser.
2. Researcher creates an XPath query that identifies the scientific properties they are looking for (following the QCDML schema for an ensemble). ILDG Browser includes a query constructor module for users who are not familiar with XPath.
3. Researcher submits query to regional grid Metadata Catalogues and waits for results.
4. Researcher browses results from query and identifies an interesting ensembles.
5. Researcher uses ILDG Browser to retrieve a list of the (LFNs for) configurations within the particular ensemble.
6. Researcher generates a proxy certificate and then uses the ILDG 'getURL' client to contact regional grid file catalogues and establish SURLs for each of the configurations.
7. Researcher uses SURLs to download configuration data to local computer, possibly using srmcp, globus-url-copy, or wget, for SRM, GSIFTP, and HTTP protocols, respectively.

   Note that, due to the size/number of the datasets and potential bandwidth constraints, this download step may take some time.
8. Researcher performs analysis on retrieved datasets.

Middleware/applications involved

- Web service containers (hosting regional grid MDCs and FCs).
- Regional grid file catalogues (e.g. gLite File Catalogue or Globus Replica Location Service).
- File transfer services/clients – SRM-compliant, GridFTP, and so on.
- Bespoke ILDG clients:
  - ILDG Browser (http://forge.nesc.ac.uk/projects/qcdgrid/).
  - ILDG getURL client (currently in development).

Notable items

- Different regional grids provide different data transfer protocols, as dictated by resource providers and historic decisions.
- At the time of writing, no common standards exist for interfacing with file catalogues from different middleware stacks (most notably, gLite File Catalogue and Globus RLS).
- At the time of writing, users have experienced poor bandwidth between Europe and Australasia. The level of network performance has, in some instances, made downloading ensembles of data impossible.