



GridPP

UK Computing for Particle Physics

Use Cases and Requirements for GridPP Portal Project

GIDON MOONT

Document History:
December 2004

High Energy Physics Group
The Blackett Laboratory
Imperial College London
South Kensington Campus
Prince Consort Road, LONDON SW7 2BW
Tel: +44 (0)207 594 7810
FAX: +44 (0)207 823 8830
email: g.moont@imperial.ac.uk
website <http://www.hep.ph.ic.ac.uk/~gidon/>

Contents

| | |
|--------------------------------------------------------------------------------------------------------|-----------|
| 1 Introduction | 3 |
| 2 Interested Parties | 3 |
| 2.1 CALICE | 3 |
| 2.2 MICE | 3 |
| 2.3 UKDMC | 4 |
| 2.4 GridPP | 4 |
| 3 Use Cases | 4 |
| 3.1 CALICE-1: Simulation | 5 |
| 3.2 CALICE-2: Storage of Raw Experimental Data | 6 |
| 3.3 CALICE-3: Batch Conversion of Raw data Formats to LCIO | 7 |
| 3.4 CALICE-4: Access and Management of Experimental and Simulated Data Stored on the Grid | 8 |
| 3.5 MICE-1: Generation of Radio Frequency Background | 9 |
| 3.6 MICE-2: Monte-Carlo Production and Digitisation of Output | 10 |
| 3.7 MICE-3: Reconstruction of Data | 11 |
| 3.8 UKDMC-1: Simulation | 12 |
| 3.9 UKDMC-2: Stripping of Raw Experimental Data | 13 |
| 3.10 UKDMC-3: Access and Management of Stripped Experimental Data Stored on the Grid | 14 |
| 3.11 D-1: Exhibition / Talk Demonstration of Portal | 15 |
| 3.12 D-2: Web available portal demonstrations | 17 |
| 4 Derived Requirements | 18 |
| 5 Conclusions | 19 |

1 Introduction

The GridPP Portal Project aims to provide easy access to the Grid through portals on a variety of platforms. In order to concentrate efforts on functionality useful to the end users, this document aims to establish use cases and derive requirements from them. The findings of this document will then effect the strategy for the subsequent development of the portal work.

2 Interested Parties

This Portal post has the aim of providing a portal framework. This framework will be based on the requirements of the users named in the proposal for this post, namely three HEP collaborations and the dissemination officer of GridPP. These four users are briefly described in this section.

2.1 CALICE

CALICE is an umbrella for several R&D projects investigating high resolution calorimetry for a linear collider. The activities are focused on a beam test to be performed in 2004-6 in which a prototype calorimeter module will be exposed to a variety of hadron and electron beams. As well as yielding information about various detector technologies, the beam test will play a crucial role in validating the Monte Carlo simulation programs which will be used to optimise the design of a full detector.

2.2 MICE

The Muon Ionisation Cooling Experiment (MICE) is a worldwide collaboration, the aim of which is to create a feasible method for cooling muons for the neutrino factory and muon collider. The experiment will be based at the Rutherford Appleton Laboratory and the UK is responsible for provision of the muon beam and surrounding infrastructure. In addition the UK physics groups are responsible for the tracker for the experiment itself. The goal is to begin to take data with one liquid-hydrogen absorber in 2006.

2.3 UKDMC

The UK Dark Matter Collaboration is a consortium of astrophysicists and particle physicists, conducting experiments with the ultimate goal of detecting rare scattering events which would occur if the Galactic dark matter consists largely of a new heavy neutral particle, a ‘neutralino’ or ‘WIMP’. Detectors are operating some 1100 m underground in a halite (NaCl) seam at the Boulby Mine of Cleveland Potash Ltd. The programme is funded by the Particle Physics and Astronomy Research Council (PPARC), and has also received generous support from Cleveland Potash Ltd.

2.4 GridPP

Part of the role of GridPP is to encourage use of Grid technologies, and to publicise how these technologies are currently working for the benefit of the UK HEP community. To facilitate this there is a dissemination officer whose roles includes encouraging individuals, experiments and businesses to start using the Grid. This Portal Project will contribute to that aspect of dissemination by providing an eye catching interface to a working grid utilisation.

3 Use Cases

The purpose of a use case is to find out, from the users of the system being built, what it is that they would like to be able to do. They should not be providing specific details of how they will do this, but should be providing information on the basic sequence of events they expect to follow.

A Use case therefore needs to describe, in broad terms, the sequence of events that the system should be able to complete for the user. It also needs to note what previous conditions need to be met, and make a note of what postconditions will equate with success. It is also useful to know who will carry out these tasks, and when. Assumptions should be made clear.

It is also recognised that various modifications to the basic sequence of events may be desired, and these should also be recorded as extended actions.

3.1 CALICE-1: Simulation

- Basic actions:** Collate the needed executable and data files into the portal. Submit Grid job of order 100,000 -> 1,000,000 Monte Carlo simulated events, split into sensible chunks. Store results on a Grid Storage Element. Both raw data and Linear Collider I/O (LCIO) outputs. Log all Grid storage changes, including user and time of action.
- Discussed with:** Dan Bowerman on 6th October 2004.
David Ward, Paul Dauncey on 13th October 2004.
- Frequency of Action:** Daily from end of 2004 until mid 2006 and possibly ongoing afterwards.
- People Involved:** A dozen + from UK and abroad.
- Preconditions:** There exists a MOKKA/GEANT executable, compiled statically to avoid AFS problems, on Enterprise Linux. The worker node on the Grid can access the mySQL geometry databases for MOKKA, or another solution for providing MOKKA with the geometry file is found. Users will somehow be able to create proxies for job submission.
- Postconditions:** Resubmission if failure rate > 5%.
< 24 hour turnaround.
- Extended actions:** Submit multiple jobs with either different GEANT models (meaning different compiled binaries) or with a single parameter changing in an iterative manner.
- Assumptions:** Users will have Certificates and a Virtual Organisation (VO) membership.

3.2 CALICE-2: Storage of Raw Experimental Data

- Basic actions:** Upload raw experimental data files to a Grid Storage Element from a testbed site.
Log all Grid storage changes, including user and time of action.
- Discussed with:** Dan Bowerman on 6th October 2004.
David Ward, Paul Dauncey on 13th October 2004.
- Frequency of Action:** During experiment runs, the first of which will be from December 2004 – Spring 2005 at DESY. There will be other periods at later dates, each lasting for several months.
- People Involved:** Those members of the collaboration at the testbed sites; DESY (Germany), FNAL (US), IHEP (Russia).
- Preconditions:** None.
- Postconditions:** The uploaded data will be accessible through the portal and using Grid software.
- Extended actions:** Convert data files to LCIO format, and store the LCIO format as well as the raw data.
- Assumptions:** Users will have Certificates and a VO membership.

3.3 CALICE-3: Batch Conversion of Raw data Formats to LCIO

- Basic actions:** Browse Grid data/meta-data catalogues.
Select Grid stored raw data. This could be from either simulation or experiment.
Convert the raw data to LCIO format.
Store the new LCIO format data on a Grid Storage Element.
Report back to user on completion.
Log all Grid storage changes, including user and time of action.
- Discussed with:** Paul Dauncey on 14th October 2004.
- Frequency of Action:** Anytime after either simulated or experimental data has been put onto the Grid.
- People Involved:** Any member of the collaboration.
- Preconditions:** A Grid compatible (compiles on correct platform) binary to do conversions.
- Postconditions:** Accessible storage on the Grid.
- Extended actions:** None.
- Assumptions:** Users will have Certificates and a VO membership.

3.4 CALICE-4: Access and Management of Experimental and Simulated Data Stored on the Grid

- Basic actions:** Browse Grid data/meta-data catalogues.
Select Grid stored data. This could be from either simulation or experiment, raw or LCIO.
Download the data - i.e. to local machine.
- Discussed with:** Dan Bowerman on 6th October 2004.
- Frequency of Action:** From early 2005 onward. Sporadic, depending on data coming from experiments and simulations.
- People Involved:** Potentially anyone in the CALICE collaboration.
- Preconditions:** That previous Use Case postconditions have been met - i.e. that both simulation and experimental data has been stored on the Grid.
- Postconditions:** Logs of all Grid storage changes, including user and time of action.
- Extended actions:** Delete data from the Grid that is of no further use. Any such deletion must be logged.
- Assumptions:** Users will have Certificates and a VO membership.

3.5 MICE-1: Generation of Radio Frequency Background

- Basic actions:** Collate the needed software - a GEANT4 based statically linked executable - along with datacard and RF spectrum files.
Submit Grid job.
Retrieve results.
- Discussed with:** Ken Long on 22nd October 2004.
Malcolm Ellis on 25th October 2004.
- Frequency of Action:** From as soon as available, with growing need for foreseeable future.
- People Involved:** Malcolm Ellis + others
- Preconditions:** GEANT4- software compiled on Enterprise Linux.
Users will somehow be able to create proxies for job submission.
- Postconditions:** Resubmission if failure rate > 5%.
- Extended actions:** The datacard files will need to contain different seeds to be used by the CLHEP random number generator. At the moment this is done by hand. Automatic selection of a new seed could be done by a MICE provided component that runs within the portal.
The job will publish its progress, allowing for the portal to provide more detailed tracking information.
- Assumptions:** Users will have Certificates and a VO membership.

3.6 MICE-2: Monte-Carlo Production and Digitisation of Output

- Basic actions:** Collate the needed files. The first step requires a GEANT4 based statically linked executable, along with datacard file and in some cases the RF Background files. The second step requires a statically linked executable to do the digitisation.
Submit Grid job. It will run as a two step process.
Retrieve results.
- Discussed with:** Ken Long on 22nd October 2004.
Malcolm Ellis on 25th October 2004.
- Frequency of Action:** From as soon as available, with growing need for foreseeable future.
- People Involved:** Malcolm Ellis + others
- Preconditions:** GEANT4- software compiled on Enterprise Linux.
For the RF files to exist, MICE-1 should successfully complete (although these files could come from another source).
Users will somehow be able to create proxies for job submission.
- Postconditions:** Resubmission if failure rate > 5%.
- Extended actions:** The datacard files will need to contain different seeds to be used by the CLHEP random number generator. At the moment this is done by hand. Automatic selection of a new seed could be done by a MICE provided component that runs within the portal.
The job will publish its progress, allowing for the portal to provide more detailed tracking information.
- Assumptions:** Users will have Certificates and a VO membership.

3.7 MICE-3: Reconstruction of Data

- Basic actions:** Collate the needed files - a Reconstruction statically linked executable - along with datacard and output from MICE-2 use case.
Submit Grid job.
Retrieve results.
- Discussed with:** Malcolm Ellis on 25th October 2004.
- Frequency of Action:** From as soon as available, with growing need for foreseeable future.
- People Involved:** Malcolm Ellis + others
- Preconditions:** Reconstruction software compiled on Enterprise Linux.
For the needed input files to exist, MICE-2 should successfully complete.
Users will somehow be able to create proxies for job submission.
- Postconditions:** Resubmission if failure rate > 5%.
- Extended actions:** The job will publish its progress, allowing for the portal to provide more detailed tracking information.
- Assumptions:** Users will have Certificates and a VO membership.

3.8 UKDMC-1: Simulation

- Basic actions:** Collate the needed executable and data files into the portal.
Submit Grid job of order 100 million Monte Carlo simulated events, split into sensible computational chunks by the portal.
Retrieve results.
- Discussed with:** Alex Howard 19th October 2004.
- Frequency of Action:** Spring 2005 onwards.
- People Involved:** A dozen + from UK (Imperial,Sheffield,RAL) and abroad.
- Preconditions:** Software (GEANT of various flavours) and required Environment variables adjusted to be able to run on Grid.
Random Seed problem understood and accounted for.
Users will somehow be able to create proxies for job submission.
- Postconditions:** Resubmission if failure rate > 5%.
Would want to know why failures occurred - Grid problem or software (random seed) problem? Try to return what data output does exist after a failure.
- Extended actions:** Submit multiple jobs with iteratively changing Macro files (that control the GEANT executable's runtime behaviour).
Use shell scripts on the portal to automate job submission.
- Assumptions:** Users will have Certificates and a VO membership.

3.9 UKDMC-2: Stripping of Raw Experimental Data

- Basic actions:** Browse Grid data/meta-data catalogues.
Select Grid stored data.
Request batch stripping of data using given software.
Save results on a Grid Storage Element.
Retrieve results.
- Discussed with:** Alex Howard 19th October 2004.
- Frequency of Action:** During experiment runs, starting in Spring 2005.
- People Involved:** Any member of the collaboration who is responsible for determining how the raw data will be stripped.
- Preconditions:** The data is being batch moved from Boulby Mine to the ATLAS Data Storage at RAL. There needs to be an interface to ATLAS from the Grid, along with the necessary catalogues of the data.
The program to do the stripping needs to be runnable on the Grid.
- Postconditions:** Accessible storage on the Grid.
Logs of all Grid storage changes, including user and time of action.
- Extended actions:** None.
- Assumptions:** Users will have Certificates and a VO membership.

3.10 UKDMC-3: Access and Management of Stripped Experimental Data Stored on the Grid

- Basic actions:** Browse Grid data/meta-data catalogues.
Select Grid stored data.
Download the data - i.e. to local machine.
- Discussed with:** Alex Howard 19th October 2004.
- Frequency of Action:** From mid 2005 onward.
- People Involved:** Any member of the collaboration who is wanting to analyse the stripped data.
- Preconditions:** That the previous UKDMC-2 Use Case postconditions have been met.
- Postconditions:** Logs of all Grid storage changes, including user and time of action.
- Extended actions:** Delete data from the Grid that is of no further use. Any such deletion must be logged.
- Assumptions:** Users will have Certificates and a VO membership.

3.11 D-1: Exhibition / Talk Demonstration of Portal

- Basic actions:** Submit a job. This job could be either a HEP Monte Carlo portal case (fewer events than normal) or a protein science job (protein docking using FTDock?) or something else.
View a job's status. Ideally using a graphical interface. This could mean using a version of the Imperial College LCG Monitoring applet.
View a job's results. As the job sections complete, build up a fuller picture of the results. Do this graphically.
- Discussed with:** Sarah Pearce 21st October 2004.
- Frequency of Action:** At conferences.
- People Involved:** Dissemination officer or management person of GridPP.
- Preconditions:** The jobs will be ready to go from the portal - i.e. all correct executables and other files will already be on portal.
Viewing the job's status could mean modification of the applet to be able to identify the specific jobs submitted from the portal. Another solution could be to run our own Resource Broker (which would be the only RB queried by the applet).
Viewing the job's results in graphical form would require running something to generate the pictures on the portal, or possibly on a remote site.
- Postconditions:** May want to somehow store (screen capture?) the process for later quick demonstrations. This is especially true of the viewing parts, as a job could run over several hours.

Extended actions: View an overall log of how the job ran - number of nodes used/SI2K/queue times/etc.

Have a parallel running program, on a single machine (possibly same as the one being used to interface to the portal - but this may cause the machine to be sluggish), also reporting on its progress. This compare and contrast would be particularly good for an exhibit.

Assumptions: Users will have Certificates and a VO membership. They will have them with them at the right time and place.

3.12 D-2: Web available portal demonstrations

- Basic actions:** The portal will continually run at least one example of the same type as D-1.
People browsing the web will be shown a view of the job's status and results, from a non-secure (non-Certificate requiring) side of the portal.
The portal will point people to further sources of information, both about the grid and about the science connected with the job.
- Discussed with:** Sarah Pearce 21st October 2004.
- Frequency of Action:** Continual once in existence.
- People Involved:** Anyone in the world.
- Preconditions:** D-1 is fulfilled.
The portal machine is able to cope with the browser load.
Appropriate copy and links are available to make the portal more than just about the computing - it should be as much about the science being done.
- Postconditions:** Web-logs being analysed.
- Extended actions:** Rather than continually running new jobs once the old one has finished, a button will become active to allow that one lucky person to click-start a new job.
- Assumptions:** The portal will have a proxy certificate for itself to use. This assumes this is acceptable with those in security area.

4 Derived Requirements

1. A mechanism for secure connection to the portal. This must use Grid Certificates to identify the user and establish the connection.
2. A method, or methods, to upload and download files (executables, datacards, results) to and from the portal to a users client machine. There will be a mechanism for identifying ownership of these files on the portal, and an appropriate permissions system will exist.
3. Virtual Organisations will need to be able to share work, so the permissions on files (executables and data) needs to include a group concept. File permissions should be based on the Grid Certificate and VO of the portal user who created it (or caused it to be created through a job), in a manner analogous to UNIX file permissions (including default mask).
4. Files that have been stored on a Grid Storage Element need to be easily accessible for both further jobs and for downloading to a user's client machine. This would involve the ability to browse metadata catalogues. This may also involve the portal having its own metadata on all files that pass through it onto the Grid, or are generated by jobs that the portal submits to the Grid.
5. Proxies of Grid Certificates will need to be generated and managed. It is important that this is as much a one click process as possible.
6. The following functionalities need to be made available, probably as fixed scenarios based on the use cases, and should not require the user to have any knowledge of the command structure of the Grid;
 - Job submission
 - Job tracking
 - Handling of Job results, both data files and STDOUT/STDERR
 - Statistics on completed jobs
7. If a job publishes in some form its progress, this should be included as part of the tracking.
8. Grid monitoring should be available through the portal.

9. Parameter scanning and similar repetitive process control of random seed generation to be able to be added by pluggable components provided by the collaborations.
10. The portal able to execute a shell script that will want to do multiple job submissions. The shell script would have to be filtered through a permissions control system to maintain security.
11. Some jobs may want to start on the completion of a previous job. This could either be done at the portal or through the use of DAGS.
12. For the dissemination portal, a method of graphically representing experimental results as they return from the Grid. This will be specific to the application chosen to be run as a demonstration.
13. The portals need to be easy to use, navigate, grow and maintain. They must provide enough help to make this so.

The following comment is not a requirement from any of the above Use Cases. However, it has been asked whether the portal should also provide more non-Grid specific functionalities. It is therefore noted that in the long term, it would be desired by some collaborations to have the Grid portal also act as a more general portal for their work, including such things as a diary, meetings documents, and other group related work. Although this is beyond the remit of this project, it should be held in consideration when designing the portal. An architecture that readily allows the addition of components with such functionalities would be an advantage.

5 Conclusions

Use Cases have been gathered from those named in the proposal for the GridPP Portal post, namely three HEP collaborations and the dissemination officer of GridPP. Based on the acquired information, requirements have been derived and will form the basis for the design of the GridPP Portal.