



# METADATA IN A GRID ERA.

Paul Millar<sup>1</sup>, Solveig Albrand<sup>2</sup>, Mòrag Burgon-Lyon<sup>1</sup>, Stephen Burke<sup>3</sup>, Carmine Cioffi<sup>4</sup>, Tony Doyle<sup>1</sup>, Tom Doherty<sup>1</sup>, Jerome Fulachier<sup>2</sup>, Randolph Herber<sup>5</sup>, Fabian Lambert<sup>2</sup>, Adam Lyon<sup>5</sup>, Gavin McCance<sup>6</sup>, Helen McGlone<sup>1</sup>, Wyatt Merritt<sup>7</sup>, Caitriana Nicholson<sup>1</sup>, Krzysztof Nienartowicz<sup>6</sup>, Richardo Rocha<sup>6</sup>, Stefan Stonjek<sup>1</sup>, Owen Syngé<sup>1</sup>, Julie Trumbo<sup>5</sup>, James Werner<sup>8</sup>


<http://www.gridpp.ac.uk/datamanagement/metadata/>

## Introduction

The UK has a strong involvement with HEP Grid computing, crystallised by the GridPP project. A continuation of GridPP (the GridPP2 project, 2004–2007) has been established to maintain a central focus for UK HEP Grid work. Since all experiments are involved with metadata, GridPP has established a new group to investigate the use of metadata within HEP Grid computing. The group aims to ensure that the metadata services are deployed as much as possible using common technologies, that they are built using existing common interfaces (wherever possible) so that products are interchangeable, and that they use common schema for problems that are common across the experiments. The GridPP posts are dedicated to working with various aspects of metadata but the metadata group is an international collaboration and includes many others, working for CERN, various LHC experiments, EGEE and further afield.

An important aspect of the group's work is to provide a forum in which all those working on metadata services can be informed of each other's work within this area. To this end, the group has various collaborative resources, monthly meetings and annual workshops.

The group also has a small cluster of computers that allows for software development and testing, in addition to testing of metadata solutions over international communication networks.



## The 13 core use-cases

The metadata group encompasses people working on metadata from different experiments and physics groups. These groups have conducted independent use-case investigations, producing documents that describe how a user should interact with a Grid, including any metadata services. These documents contain a wealth of information, but not targeted at metadata specifically. Moreover, the information from the different sources is often complementary.

One of the first tasks undertaken by the group was reviewing the collection of general use-case documents, identifying the corresponding requirements on any metadata service and consolidate these as metadata use-cases.

This work is available as the paper "Unlucky for Some—the Thirteen Core Use Cases of HEP Metadata."

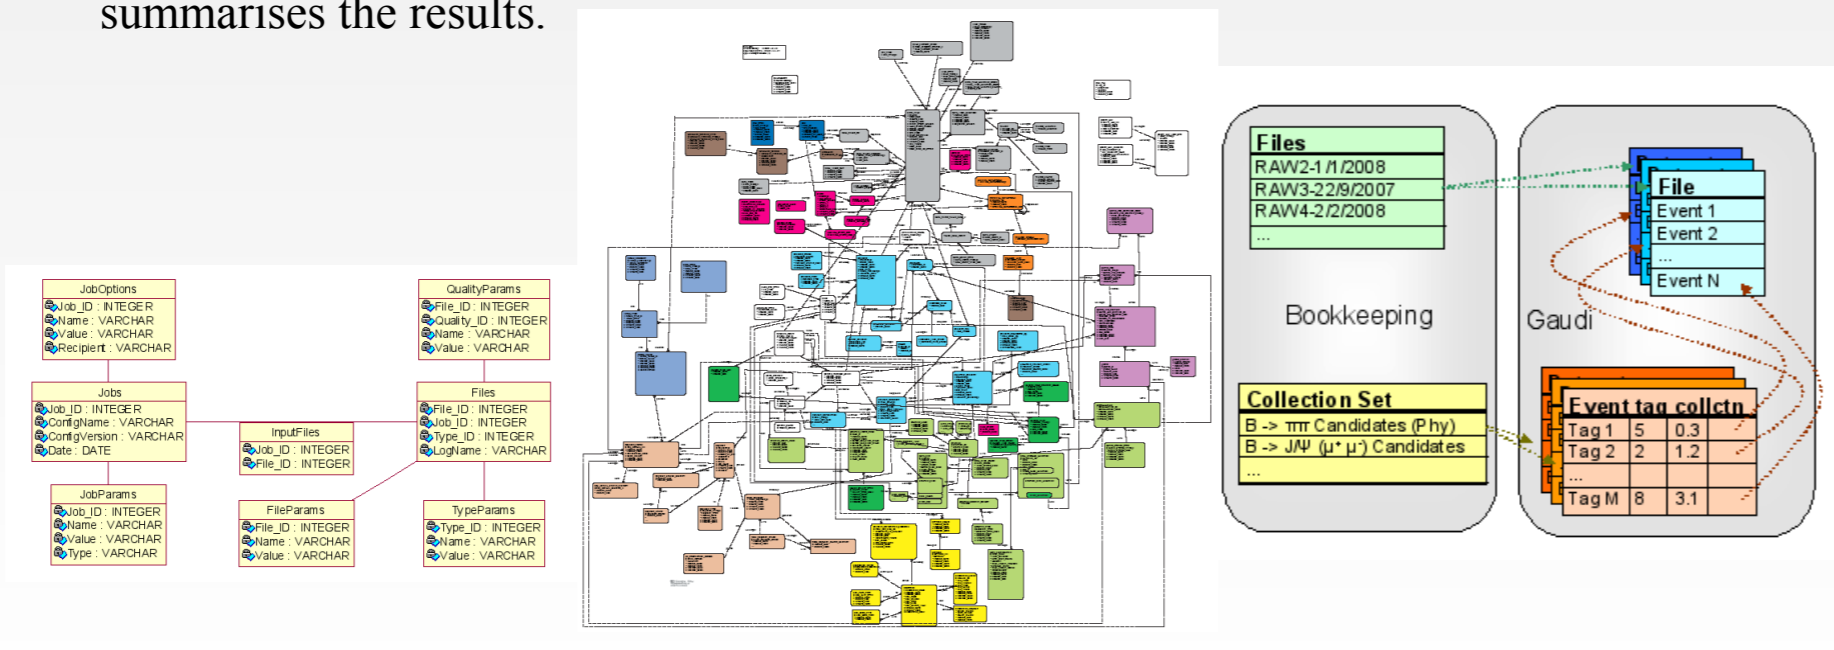
Core Use Case	Source Use Cases			
	HEPCAL	CDF	ATLAS	BaBar
1.1 Specify a new Dataset	dsreg	4.2	UC1.1, UC2.1, UC2.4, UC3.1, UC4.1, UC4.2, UC4.4	
1.2 Read Metadata for Datasets	dsmdisc	4.1	UC1.5, UC3.3	U1.5
1.3 Update Metadata for a Dataset	dsmdupd	Not Allowed	UC3.3	
1.4 Resolve Physical Data	dsaccess, remdsacc	4.6, 4.7, 4.8	UC3.3	U1.6
1.5 Download a Dataset to Local Disk	dsdownload	4.5, 4.9		
2.1 Run a Physics Simulation Program	simjob	4.4, 4.10, 4.11		
2.2 Select a Subset of a Dataset	analysis1	4.7	UC4.2, UC4.4	U1.5
2.3 Run an Algorithm over an Input Dataset	dtrans, analysis1	4.6, 4.9	UC1.4, UC1.5, UC2.2, UC2.3, UC2.4, UC3.1	U1
3.1 Submit a Job to the Grid	jobsub	4.3		U1.9
3.2 Retrieve/Access Output of a Job	jobsubaccess	4.9		U1.14
3.3 Estimate the Resource cost of Running a Job	jobestat	MW9		
3.4 Monitor the Progress of a Job	jobcont, jobmon	MW9		U2
3.5 Repeat a Previous Job	See REPCAL II	4.8		U4, U5

## Schema summary

Metadata (e.g. file locations on the Grid or event data summarising occurring Physicals) must be stored. Current consensus is to store this data within a conventional relational database, with an additional wrapper layer to allow more abstract metadata concepts.

The choice of schema for storing metadata depends on the functionality needed of the metadata service. Different projects have undertaken to store the data within different schemata, yet there is merit in understanding the benefits of these different approaches.

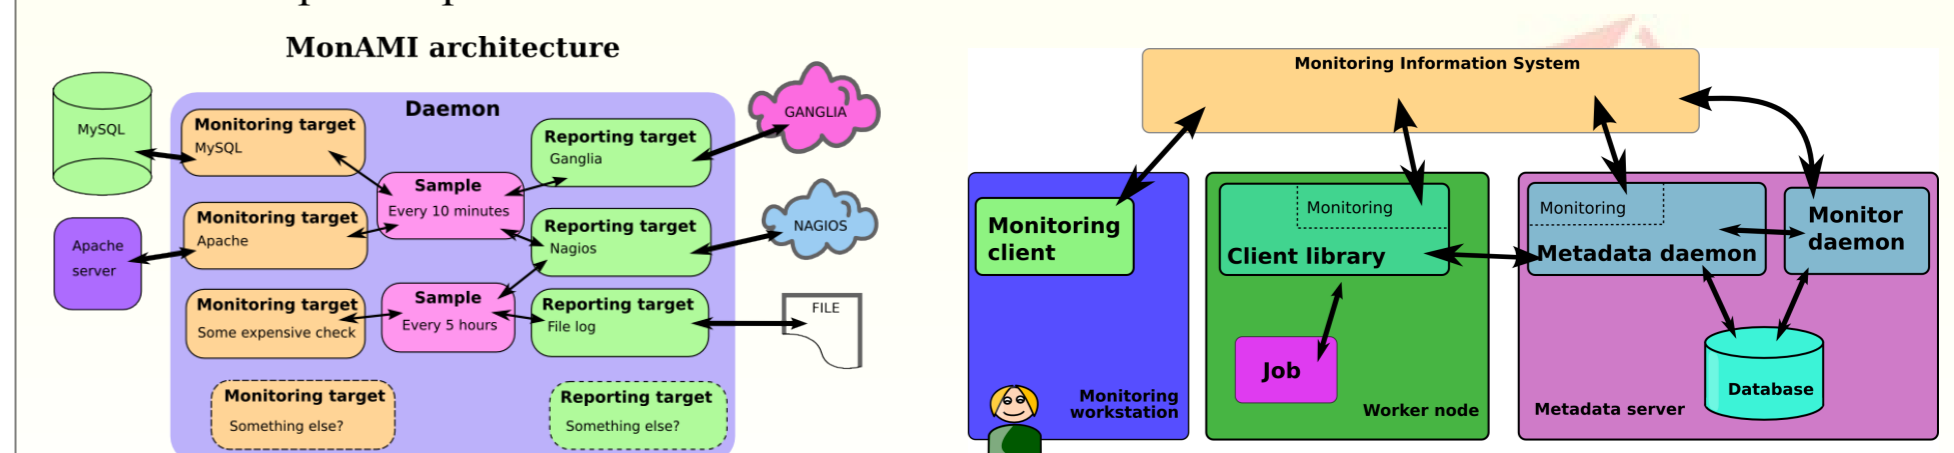
The metadata group reviewed the different schemata and produced a paper that summarises the results.



## Monitoring

Monitoring in a Grid environment is important. Yet different aspects of monitoring are useful to different people. System administrators need to know quickly if a problem is present. Virtual Organisation managers might need to observe trends in hardware usage to predict at what point will they need to upgrade the available hardware, end users might require information on their jobs' use of available resources and if their jobs were unduly slowed by external services.

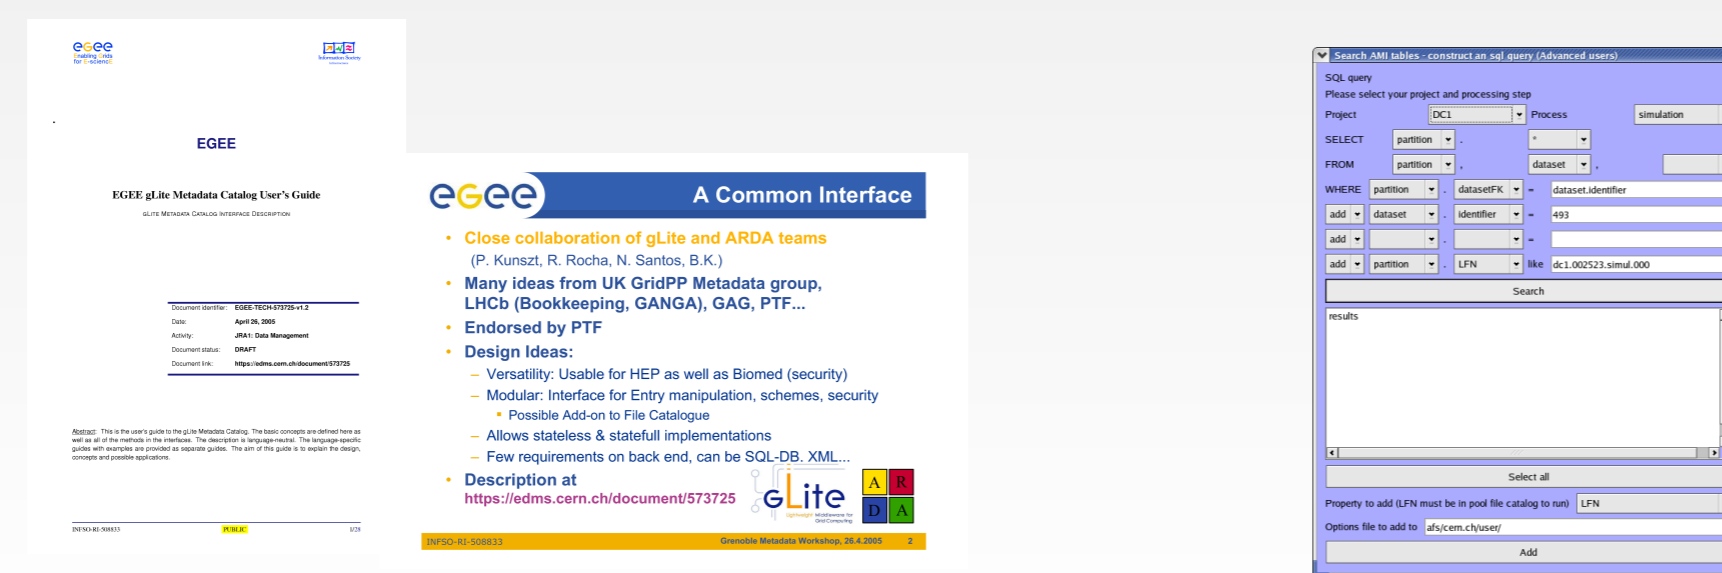
The group has produced a requirements document for such monitoring. It is also helping implement a service monitoring system that allows easy delivery of monitoring information to multiple recipients.



## Common metadata interface

One of the stated aims of the metadata group is to foster the use of standard interfaces and technologies. This allows for different technologies to incorporate the best features of other solutions. An example of this is the interface the metadata service provides.

In collaboration with LCG and EGEE, the metadata group helped establish a common interface for querying a metadata service. This interface has been adopted by several metadata servers and several metadata servers now support this interface.



## Future work

GridPP's metadata management work centres on the development of robust, scalable, secure metadata management systems capable of meeting the metadata needs of a wide variety of organisations and groups.

The focus of the work is currently to develop mechanisms for handling experiment-specific file collections and implementing the associated grid security mechanisms. Components for the distributed data management system of ATLAS, CMS and LHCb are being developed. These incorporate automated mechanisms to label file collections, gather database schema via a single interface, and define the database structures that enable individual events to be tagged for end users.

Further work on database schema evolution will allow people to work without direct connection to a central database. Later, when reconnected to the central database, their work will be submitted for automatic inclusion.

The work will continue on monitoring, in order to enable the metadata systems currently being developed to become more robust. Ultimately, this will allow end-users to discover if the metadata system has been a noticeable bottleneck for their work and if so, which component.