



**GridPP**

UK Computing for Particle Physics

# Metadata

## Technical Development Report

14<sup>th</sup> September 2004

Metadata

Steven Hanlon – University of Glasgow

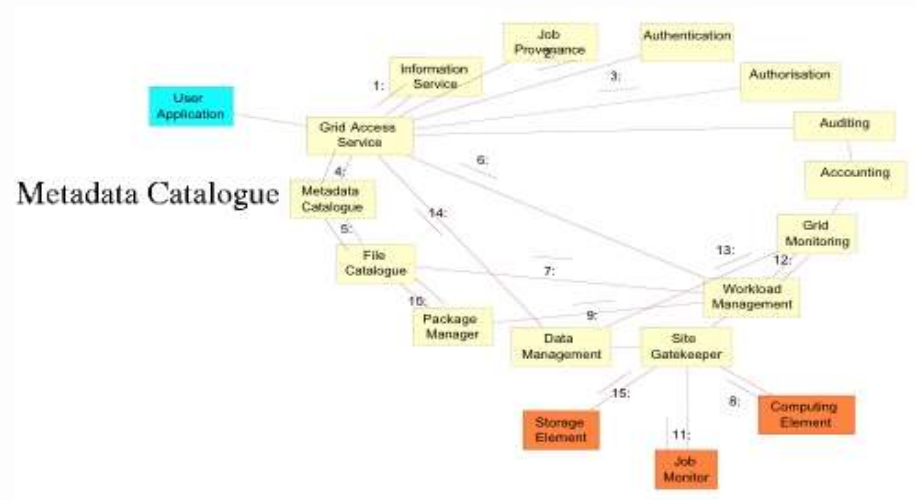


- What is Metadata?
- The Metadata Group
- Use cases and requirements
- Building Metadata Services
- Existing Metadata Services: SAM and AMI
- Conclusions
- Getting involved!



# What is Metadata?

- Lot's of answers... It depends who you are.
  - Broadly 'Data about data'
- For us, metadata is 'data about experiments' data'.
- ARDA Metadata Catalogue.
- SAM Dimensions.
- High-level information on Files, Datasets, Events.
- **Not** Sites, Jobs, File replicas etc.





# Level 1 Deliverables

Number	Title	Due Date	Status
2.1.1	Review and assess HEP metadata use cases	01-Dec-04	Not Started
2.1.2	Documenting schema for Metadata	01-Mar-05	Not Started
2.1.3	Definition of Metadata Service	01-May-05	Not Started
2.1.4	Implementation of appropriate Metadata service	On going	OK
2.1.5	Common monitoring requirements	01-Aug-05	Not Started
2.1.6	Assisting wider adoption of Metadata service	On going	OK
2.1.7	Review of Metadata service definition	01-Jan-06	Not Started
2.1.8	Throughput testing of database access via Metadata	01-Apr-06	Not Started
2.1.9	Contribution to EGEE developments	On going	OK
2.1.10	Documentation on distributed databases	01-Mar-07	Not Started
2.1.11	Optimisation of data access via metadata	01-Sep-06	Not Started
2.1.12	Final evaluation of Metadata components	31-Aug-07	Not Started

- 12 high level deliverables in GridPP2 ProjectMap
- Ongoing...
- (2.1.1) use case review document: v0.1
- (2.1.6) start-up meeting followed by monthly general meetings, wiki, web pages..



# Metadata Group Aims

- Examine commonalities of metadata handling in HEP.
  - At technology, interface and schema levels.
- Help deploy metadata services using common, standard technologies.
- Use standard DB schema for common problems.
- Work together with EGEE and LCG projects.



# Who is Involved?

- Representatives of many experiments:
  - Grenoble AMI team from ATLAS;
  - Fermilab/Glasgow SAM team from CDF/D0;
  - Oxford LHCb metadata;
  - Bristol CMS metadata;
  - Manchester BaBar metadata.
- Plus GridPP, LCG and EGEE generic metadata representatives from Glasgow, RAL and CERN.



# Areas of Interest

- Have identified 7 areas of interest:
  - Use Cases - getting a handle on requirements.
  - Deployment architecture - inside the black box.
  - Services - how to present the metadata service to the outside world?
  - Query Languages and Interfaces - how does the user communicate with the metadata DB?
  - Tools - what's out there that we can use?
  - Monitoring - maintenance, file access patterns...
  - Security



- First step is to identify:
  - Common requirements for metadata services;
  - Potential for common approaches; and
  - Existing ideas that can be shared.
- Have carried out a survey of all known HEP use case documents.
  - HEPCAL, CDF, ATLAS ADA, BaBar...
  - Identify use cases that are fundamental to a GRID system and have metadata implications.
  - Have identified 13 common use cases from a list of ~80.



# A Common Use Case

- For example...

## **Select a subset of a dataset.**

Core analysis activity. The user constructs a metadata query and it is applied to the input datasets. A new dataset is produced consisting of all data satisfying the query. The query may act on the input dataset's constituent datasets, files, events or some mixture. The copy may be shallow (e.g. output a tag dataset) or deep, for more efficient access.

## **Metadata Implications**

It must be possible to break a query down into dataset, file and event levels, and apply its elements on the appropriate databases. It must be possible to create a dataset that exists purely as metadata (e.g. a tag collection.) It must be possible to record the relationship between the input and output datasets.

## **Source Use Cases**

HEPCAL: analysis1

CDF: 4.7

ADA: UC4.2, UC4.4



- Can Metadata be changed after writing?
- CDF/D0 SAM has static metadata.
  - Small problems e.g. Dataset names reserved by flawed datasets.
  - But system is perfectly sound.
- Mutable metadata hard to manage.
  - Provenance requirements - need to recover the state of metadata items at any particular time.
  - But don't need full roll-back of database
- So, should record all updates, when they happened, who performed them and why.



# Roles and Permissions

- To what extent does the system differentiate between users?
- If metadata can be updated, who can do so?
- SAM and AMI both work on a 'flat' user permissions system at present.
  - Everyone equal in the eyes of the system.
  - Not really suitable if updating is possible.
- Existing Use Cases don't define this. Proposal:
  - Admins - read/write.
  - Project Admins - read + limited write.
  - Users - read and write new only.



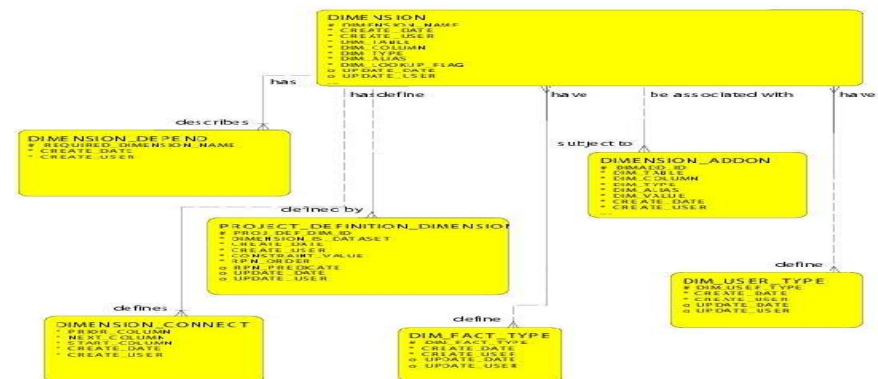
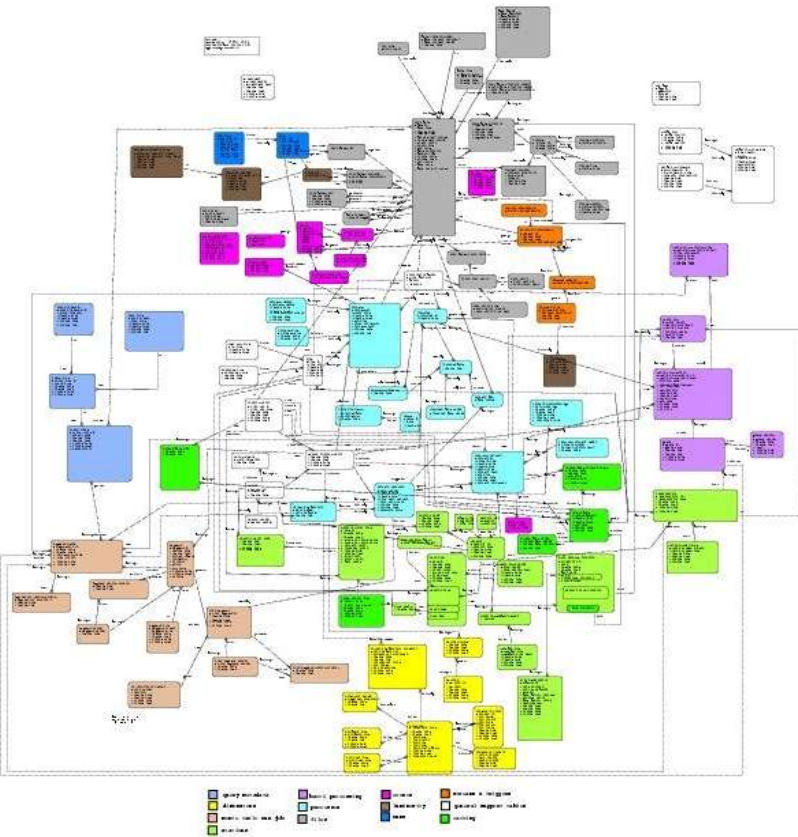
# Building Metadata Services

- One database (SAM) vs. Many Databases (ARDA)
  - Replicas
  - Separation into distinct catalogues
- What databases have useful features?
  - Consistency checking (business logic)?
  - Distributed?
  - XML based?
- Useful schema designs?
  - Keyword-Value pairs for flexibility (See later).
  - Self-description for generic interfacing (used in AMI).

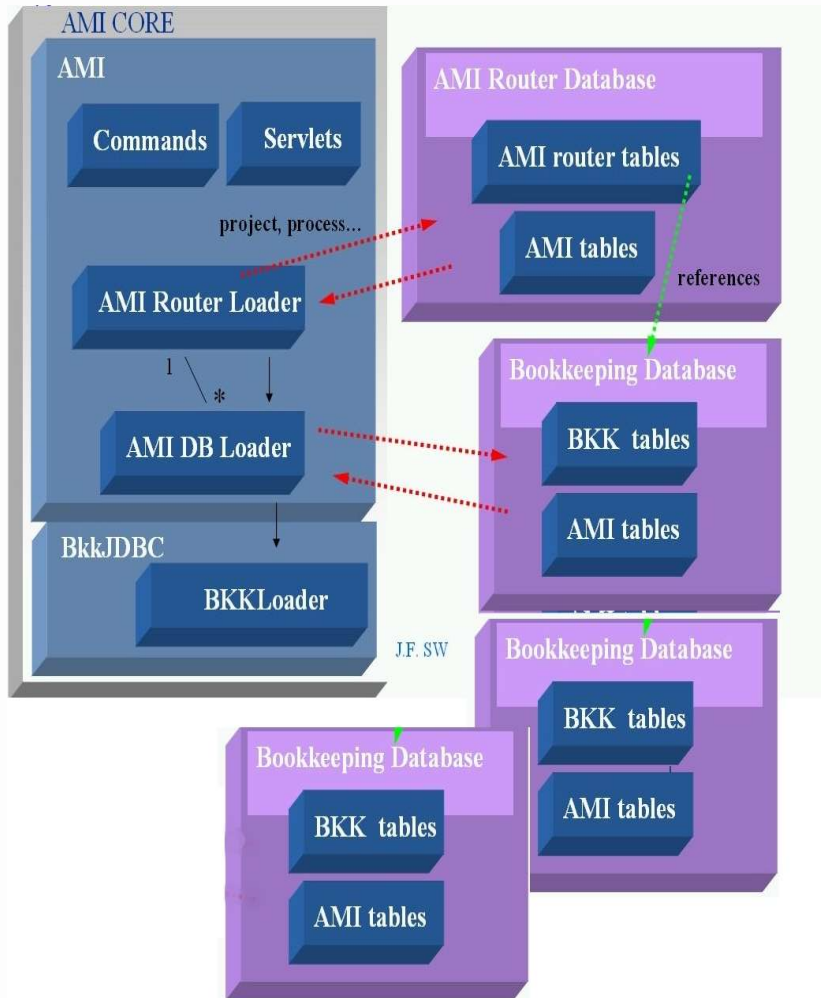


# The SAM Metadata Service

- Mature metadata service for running experiments.
- Uses single, central Oracle database for all metadata.
- Flexible extension via Dimensions (yellow).



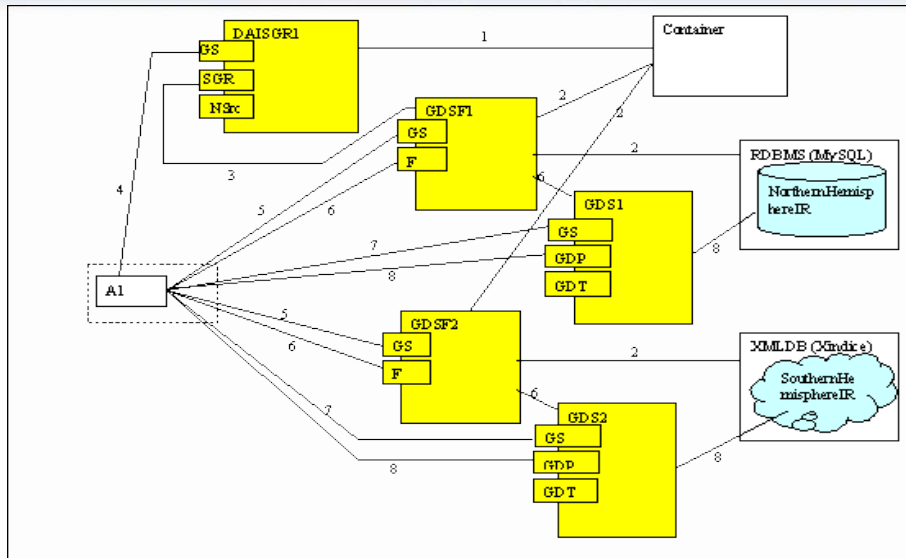
- Communication via CORBA.



- Flexible system that allows databases to be built for individual applications.
- Data can be split across several databases via router databases.
- Generic Web Service interface on top of self describing databases.
- Runs on MySQL, but should work on any DB.



# OGSA-DAI



- Allows access to data sources inside Grid standards framework.
- Data sources wrapped in Grid Services.
  - Services are registered with a Registry service...
  - ...which provides a service to access his chosen data source.
- Vision - app-independent middleware.

OGSA-DAI Distributed Query

OGSA-DAI Basic Services

Data Grid Infrastructure – Location, Delivery, Replication...

Resource Grid Infrastructure – OGSA...

Database, Communication, OS... Technology



# Keyword-Value Pairs

- Extremely flexible way to handle metadata.
- Uses column of keywords and column of values.
  - Search for 'keyword' = 'value'

<b><u>Keyword</u></b>	<b><u>Value</u></b>
Transverse Energy Cut	25 GeV
Date Created	10/02/07
Created by	Amy
Monte Carlo Program	ARIADNE
Card File	Amy.card
<i>User Keyword</i>	<i>User Value</i>
...	...

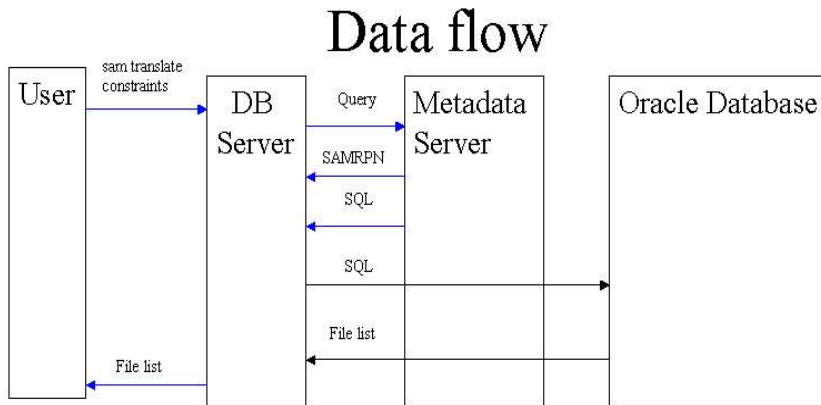
- But...
  - Slow to search - need to perform 2 matches.
- Improved method:
  - Identify frequent types of query.
  - Build proper tables accordingly.
  - How often do we need to do this?



- To what extent should the metadata system be distributed to maximise user access and efficiency?
  - Distribution removes single points of failure and can decrease response time.
  - SAM is centralised and has excellent up-time...
    - But do they just have great people?
    - Up-time vs. user access.
- What about offline 'snapshot' replicas of databases?
- How often does a replicated database change?
- Can we use databases with built-in distribution?



# Interfaces - SAM



- Don't want users trying to write SQL!
  - But databases understand SQL...
- SAM Metadata Server designed to solve this.
  - User constructs list of constraints.
  - Metadata translates to SQL.
  - CORBA-based.

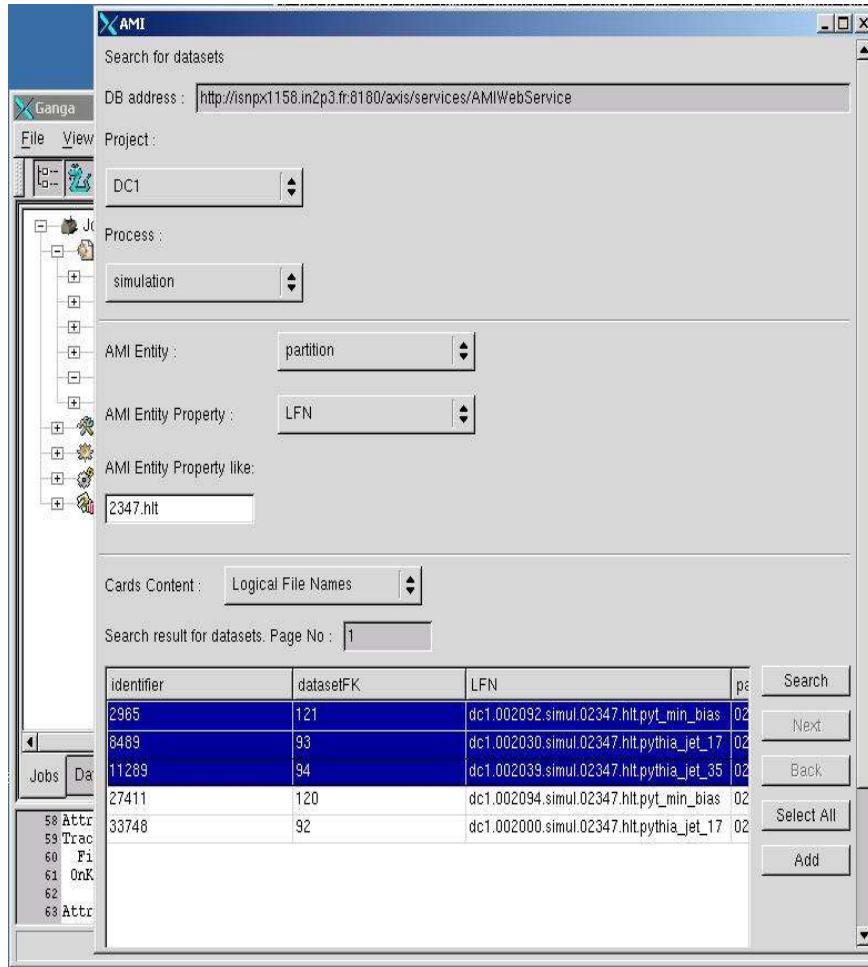
## • Examples:

```
sam translate constraints  
--dim="cdf.dataset  
stink2"
```

```
sam translate constraints  
--dim="GLOBAL.REQUESTID  
27"
```



# Interfaces - AMI



- AMI uses a Web Service Interface (SOAP).
- Very simple to create thin clients.
- Includes 'GetCommands'.
- PyAMI by Alvin Tan is a Python client for AMI.
- Has been used in a prototype GANGA interface to AMI by Chris Collins-Tooth.



- Lots to think about - this talk has been mostly questions...
- However:
  - Common metadata use cases have been identified in a survey of all known HEP use case documents.
  - Areas of common interest have been discovered...
  - ... and a cross-experiment/project team assembled.
  - Cross-project developments getting under way:
    - GANGA - AMI interface.
    - Keyword-Value pair development from AMI and SAM.
    - ...



- Lots of work and thinking to be done.
- Web Page:  
<http://www.gridpp.ac.uk/datamanagement/metadata/>
- You can see our work in progress on our Wiki:  
<http://www.astro.gla.ac.uk/users/paulm/metadata/>
- Mailing list: [metadata@physics.gla.ac.uk](mailto:metadata@physics.gla.ac.uk)
- Feel free to check what we're up to and join in!