AMGA
Metadata Access on the Grid

Mike Mineter

Acknowledgements

- **This presentation primarily consists in slides from:**
  - Tony Calanducci
    - Third EELA Tutorial for Managers and Users
    - Rio de Janeiro, 26-30 June 2006
  - Nuno Santos, Birger Koblitz
    - 20 June 2006
    - Workshop on Next-Generation Distributed Data Management
  - Patricia Méndez Lorenzo: UNOSAT application using AMGA
    - User Forum
    - CERN, 1st March 2006
    - http://indico.cern.ch/materialDisplay.py?contribId=23&sessionId=11&materialId=slides&confId=286

Contents

- Background and Motivation for AMGA
- Examples
- Interface, Architecture and Implementation
- Metadata Replication on AMGA
- gLibrary
- Further information

Metadata on the GRID

- **Metadata is data about data**
- **On the Grid: information about files**
  - Describe files
  - Locate files based on their contents
- **But also simplified DB access on the Grid**
  - Many Grid applications need structured data
  - Many applications require only simple schemas
    - Can be modelled as metadata
  - Main advantage: better integration with the Grid environment
    - Metadata Service is a Grid component
    - **Grid security**
    - Hide DB heterogeneity

AMGA Implementation

- **AMGA – ARDA Metadata Grid Application**
  - ARDA: A Realisation of Distributed Analysis for LHC
- **Now part of gLite middleware**
  - Official Metadata Service for EGEE
  - Also available as standalone component
- **Expanding user community**
  - HEP, Biomed, UNOSAT…
  - More on this later

Metadata Concepts

- **Some Concepts**
  - Metadata - List of attributes associated with entries
  - Attribute – key/value pair with type information
    - Type – The type (int, float, string,…)
    - Name/Key – The name of the attribute
    - Value - Value of an entry's attribute
  - Schema – A set of attributes
  - Collection – A set of entries associated with a schema
  - Think of schemas as tables, attributes as columns, entries as rows

Examples

- **LHCb-bookkeeping**
  - Migrated bookkeeping metadata to ARDA prototype
    - 20M entries, 15 GB
    - Large amount of static metadata
  - Feedback valuable in improving interface and fixing bugs
  - AMGA showing good scalability
- **Ganga**
  - Job management system
    - Developed jointly by Atlas and LHCb
  - Uses AMGA for storing information about job status
    - Small amount of highly dynamic metadata

Biomed

- **Medical Data Manager – MDM**
  - Store and access medical images and associated metadata on the Grid
  - Built on top of gLite 1.5 data management system
  - Demonstrated at last EGEE conference (October 05, Pisa)
- **Strong security requirements**
  - Patient data is sensitive
  - Data must be encrypted
  - Metadata access must be restricted to authorized users
- **AMGA used as metadata server**
  - Demonstrates authentication and encrypted access
  - Used as a simplified DB

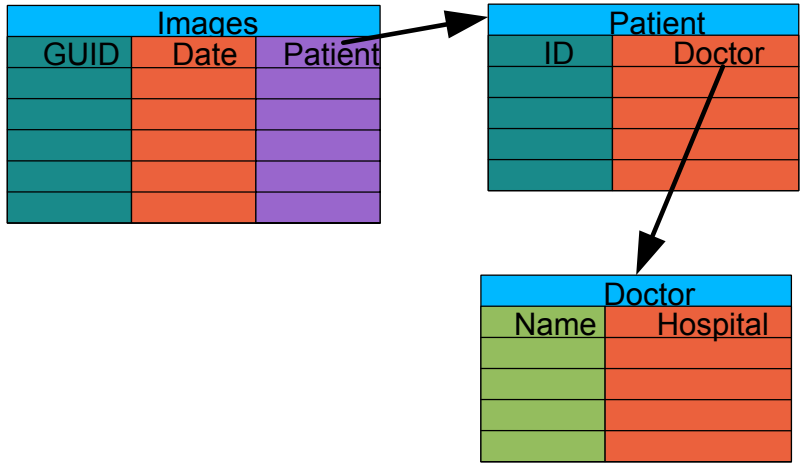

- **More details at**
  - https://uimon.cern.ch/twiki/bin/view/EGEE/DMEncryptedStorage

UNOSAT Presentation

UNOSAT is a United Nations Initiative

Objectives

- Provide the humanitarian community with access to satellite imagery and Geographic Information System services
  - Reduce disasters and plan sustainable development
- Ensure cost-effective and timely products

Core Services

- Humanitarian Mapping
- Image Processing

VEGETATION – 1 Km

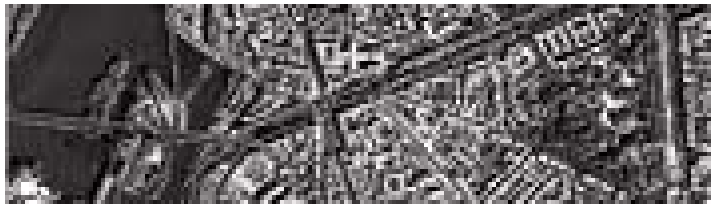

IKONOS – 1m

One step further: GRID

- Potential Bottlenecks:
  - UNOSAT beginning to suffer from limited capacity and processing power
  - Multiple satellites being launched
  - Larger and larger storage capacity needed
- In summer 2005 we have provided a whole structure at CERN for UNOSAT
  - UNOSAT Virtual Organization (VO)
  - 3.5TB in CASTOR
  - Computing Elements, Resource Brokers
  - Collaboration with ARDA group
  - AFS area of 5GB

We have provided the whole GRID infrastructure At CERN

- We have run some UNOSAT tests (images compression) inside the GRID environment (quite successful)
- The framework developed for in principle for Geant4 (See Alberto Ribon's presentation [49]) has been adapted for UNOSAT needs

UNOSAT Production inside GRID

- UNOSAT provided us with a set of images for testing
- Associated to each image a metadata file was included
  File name, directory path, geographical coordinates
- Steps:

STORAGE LEVEL

- Copy and registration of the images in Castor@CERN
  - Use of the LFC Catalog
- Parse the metadata files to extract the different metadata
- Use of the AMGA tool to parse metadata to location of the files

COMPUTING LEVEL

- Use of compression tools to compress images inside LCG resources
- Use of the general submission tool adapted to UNOSAT needs

A GRID Metadata Catalogue

LFC Catalogue

- Mapping of LFN to PFN

UNOSAT requires

- User will give as input data certain coordinates
- As output, want the PFN for downloading

The ARDA Group assists us setting up the AMGA tool for UNOSAT

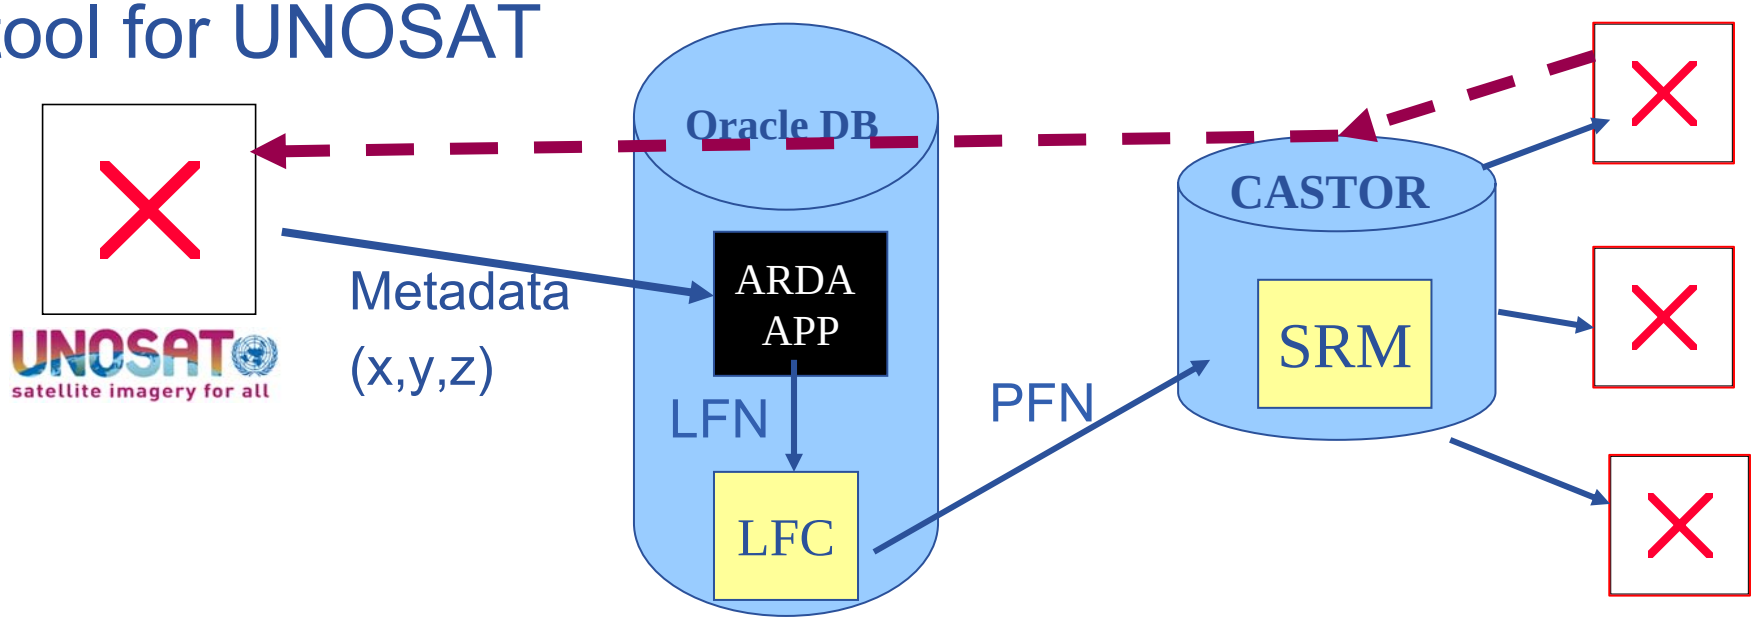

AMGA Implementation

- **AMGA Implementation:**
  - SOAP and Text frontends
  - Streamed Bulk Operations
  - Supports single calls, sessions & connections
  - SSL security with grid certs (negociated by client)
  - Own User & Group management + VOMS
  - PostgreSQL, Oracle, MySQL, SQLite backends
  - Works alongside LFC
  - C++, Java, Perl, Python clients

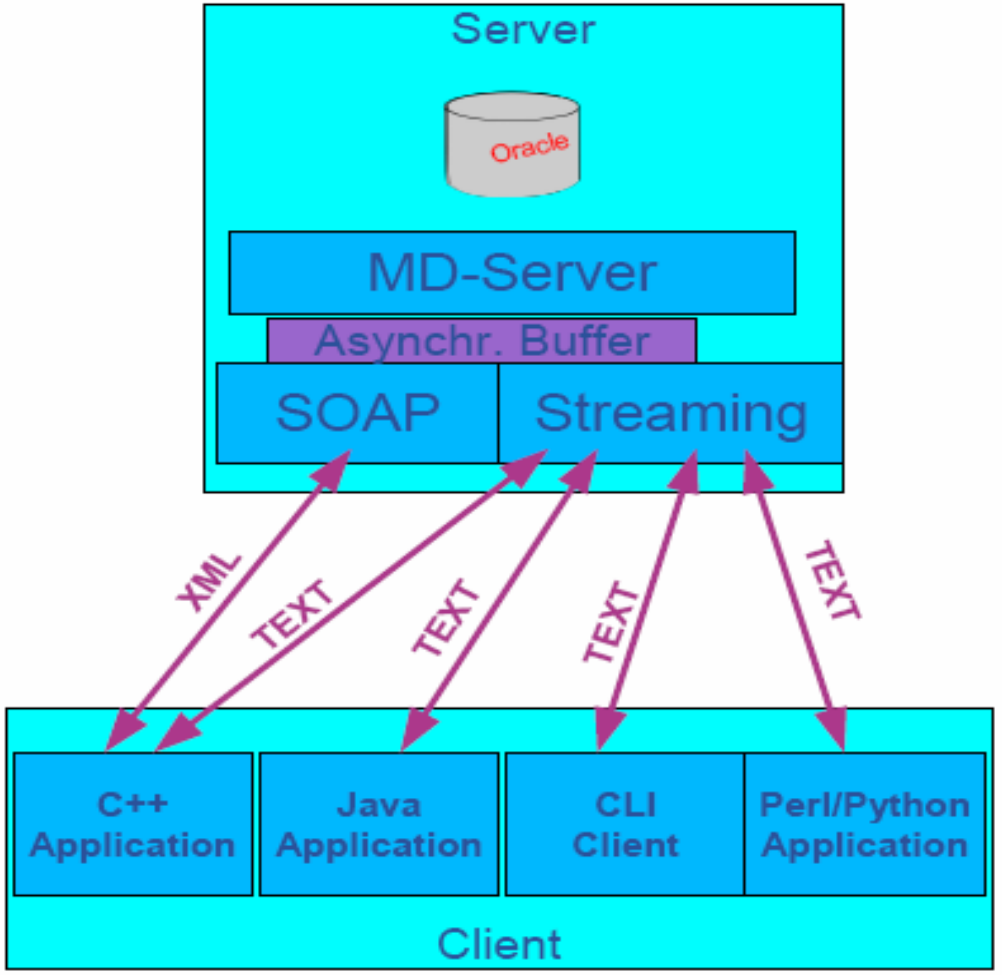

AMGA Features

- **Dynamic Schemas**
  - Schemas can be modified at runtime by client
    - Create, delete schemas
    - Add, remove attributes
- **Metadata organised as an hierarchy**
  - Schemas can contain sub-schemas
  - Analogy to file system:
    - Schema ⇔ Directory; Entry ⇔ File
- **Flexible Queries**
  - SQL-like query language
  - Joins between schemas

Security

- **Unix style permissions**
- **ACLs – Per-collection or per-entry.**
- **Secure connections – SSL**
- **Client Authentication based on**
  - Username/password
  - General X509 certificates
  - Grid-proxy certificates
- **Access control via a Virtual Organization Management System (VOMS):**

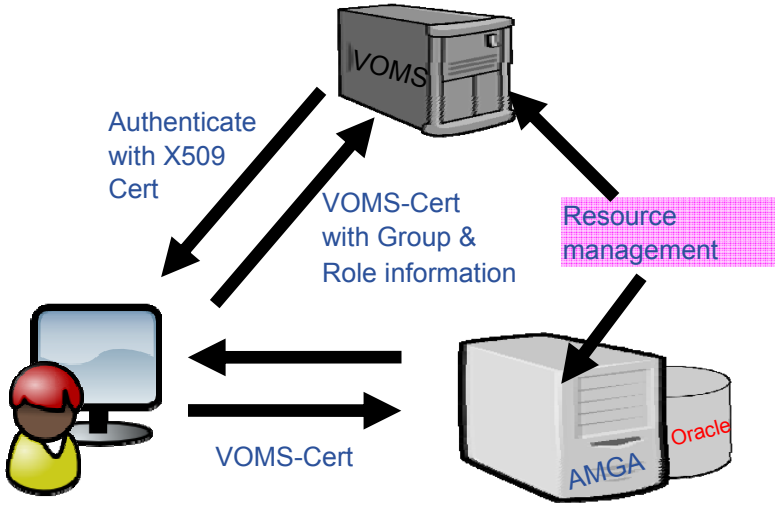

Metadata Replication

- **Currently working on replication/federation mechanisms for AMGA**
- **Motivation**
  - Scalability – Support hundreds/thousands of concurrent users
  - Geographical distribution – Hide network latency
  - Reliability – No single point of failure
  - DB Independent replication – Heterogeneous DB systems
  - Disconnected computing – Off-line access (laptops)
- **Architecture**
  - Asynchronous replication
  - Master-slave – Writes only allowed on the master
  - Replication at the application level
    - Replicate Metadata commands, not SQL → DB independence
  - Partial replication – supports replication of only sub-trees of the metadata hierarchy

Metadata Replication

Some use cases

Full replication

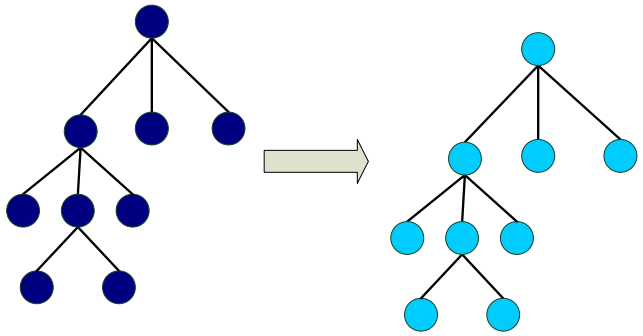

Partial replication

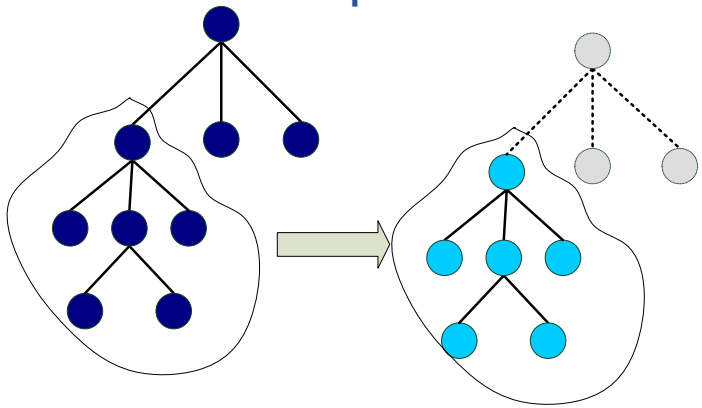

Federation

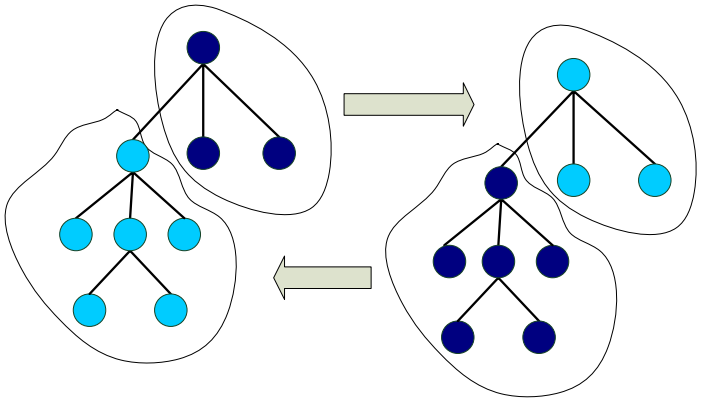

Proxy

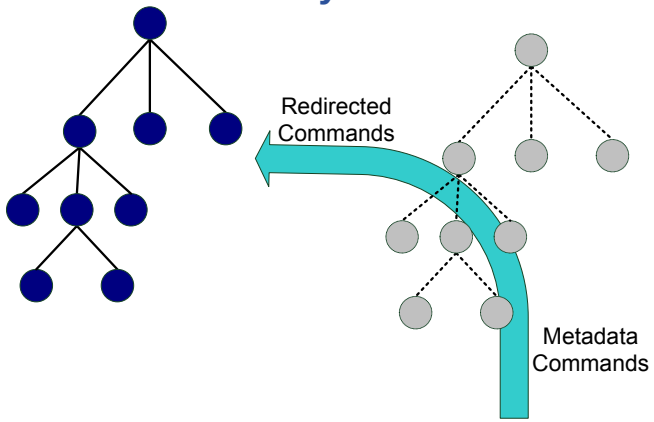

gLibrary prototype implementation

- **Files are saved on SEs and registered into file catalogues (LFC and/or FiReMan)**
- **The AMGA Metadata Catalogue is used to archive and organize metadata and to answer users' queries.**
- **gLibrary is built using the following AMGA collections:**
  - /gLibrary contains generic metadata for each entry
  - /gLAudio, /gLImage, /gLVideo, /gLPPT, /EGEEPPT, /gLDoc, … are examples of collections of "additional features" (shown later)
  - /gLTypes
    - keeps the associations between document types and the names of the collection that contains the "additional features"
    - is used by gLibrary to find out where it has to look when new document types are added into the system (extensibility)
  - /gLKeys is used to store Decryption Keys

Example of entries

Collection	/gLibrary			
Entry Names	**Attributes**			
	FileName	**PathName**	**Type**	**Submitter**
4ffaffc8-26e7-4826-b460-3d5bf08081a4	DedicatoAte.mp3	/grid/gilda/calanducci	Audio	Tony Calanducci
00454dca-a269-4b93-8a45-c4012af05600	ardizzonelarocca_is_231005.ppt.gpg	/grid/gilda/calanducci/ EGEE	EGEEDOC	Tony Calanducci

/gLibrary (continuum)				
Attributes				
SubmissionDate	**Encryption**	**Description**	**Keywords**	**CreationDate**
2006-01-05 00:00:00	false	Canzone delle vibrazioni che ha ricevuto un enorme successo tra i teenagers nel 2003	Vibrazioni	2004-02-05 00:00:00
2005-01-05 16:44:22	true	gLite Information System	R-GMA, RGMA, BDII, IS	2005-10-05 23:40

Example of gLibrary collections

Enabling Grids for E-sciencE

Collection	/gLTypes
Entry names	Attributes
	Path *(refers to a collection)*
Audio	/gLAudio
Image	/gLImage
Video	/gLVideo
Documents	/gLDOC
PowerPoint	/gLPPT
EGEEDOC	/EGEEPPT

Collection	/gLKeys
Entry names	Attributes
	Passphrase
00454dca-a269-4b93-8a45-c4012af05600	ardizzo

"additional features"

Collection	/EGEEPPT							
Entry names	Attributes							
	Title	Runtime	Author	Type	Date	Event	Speaker	Topic
00454dca-a269-4b93-8a45-c4012af05600	Information Systems	00:30:00	Valeria Ardizzione, Giuseppe La Rocca	Theorical	2005-10-23	4th EGEE Conference	Giuseppe La Rocca, Valeria Ardizzone	R-GMA, BDII

Collection	/gLAudio					
Entry names	Attributes					
	SongTitle	Duration	Album	Genre	Singer	Format
4ffaffc8-26e7-4826-b460-3d5bf08081a4	Dedicato A Te	00:03:27	Dedicato A Te	Pop	Le Vibrazioni	MP3

gLibrary Security

- **User Requirements:**
  - a valid proxy with VOMS extensions
  - VOMS Role and Group needed to be recognized by gLibrary as a contents manager.
- **3 kinds of users:**
  - gLibraryManager: (s)he can create new content type and allows a generic VO user to become gLibrarySubmitter
  - gLibrarySubmitters: they can add new entries and define access rights on the entries they create.
    - Fine-grained permission (reading, writing, listing, decrypting) settings on each entry: whole VO members, VO groups, list of DNs
  - generic VO users: browse and make queries (on entries they have access to)
- **Basic level of cryptography:**
  - New files saved on SEs can be encrypted beforehand with a symmetric passphrase that will be saved in /gLKeys. Only selected users (that have a specific DN in the subject of their VOMS proxy) can access the passphrase and decrypt the file.

Example: gLibrary queries

Initialize your VOMS proxy asking to be member of the gilda VO

Edit your .mdclient.config setting Login=NULL (user will be retrieved from your proxy extensions)

Log into AMGA

```
$ voms-proxy-init --voms gilda
$ voms-proxy-info -fqan
/gilda/Role=NULL/Capability=NULL
$ grep Login .mdclient.config
Login = NULL
```

Suppose we want to look for all contents about VOMS

```
Query> whoami
>> gilda
Query> selectattr /gLibrary:FILE /gLibrary:FileName /gLibrary:Type
'like(/gLibrary:Keywords, "%VOMS%")'
>> 1f6e9ac6-5c86-4599-b03b-560e0e7ea38a
>> VOMS_server_Installation.ppt.gpg
>> EGEEDOC
```

Now let's find out in which collection EGEEDOC attributes are stored

```
Query> getattr /gLTypes/EGEEDOC Path
>> EGEEDOC
>> /EGEEPPT
```

Example: gLibrary queries (II)

Now we can make a JOIN between the 2 tables to extract all the information we like

```
Query> selectattr /gLibrary:FILE /gLibrary:FileName /gLibrary:Description
/EGEEPPT:Author /EGEEPPT:Title /EGEEPPT:Event '/gLibrary:FILE=/EGEEPPT:FILE and
like(/gLibrary:Keywords, "%VOMS%")'
>> 1f6e9ac6-5c86-4599-b03b-560e0e7ea38a
>> VOMS_server_Installation.ppt.gpg
>> VOMS Server installation tutorial done in Venezuela
>> ziggy, Giorgio
>> Installing a gLite VOMS Server
>> First Latin American Workshop for Grid Administrators
```

Let's see where the passphrase to decrypt the file is stored

```
Query> selectattr /gLibrary:FILE DecryptKeyDir 'FILE="1f6e9ac6-5c86-4599-b03b-
560e0e7ea38a"'
>> 1f6e9ac6-5c86-4599-b03b-560e0e7ea38a
>> /DLKeys/gildateam
```

But …

```
Query> getattr /gLKeys/gildateam/1f6e9ac6-5c86-4599-b03b-560e0e7ea38a
Passphrase
Error 4: Permission denied
```

Because gilda is not a member of the gildateam group

Further information

- **on AMGA and gLibrary:**
  - http://indico.eu-eela.org/conferenceTimeTable.py?confId=37
  - (go to day 3 for the AMGA tutorial)

- **AMGA Web Site**

  `http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/`

Conclusion

- **AMGA – Metadata Service of gLite**
  - Useful for simplified DB access
  - Integrated in the Grid environment (Security)
- **Replication/Federation under development**
- **Tests show good performance/scalability**
- **Already deployed by several Grid Applications**
  - LHCb, ATLAS, Biomed, …
  - DLibrary