Architecture of the gLite Metadata Service (AMGA)

Tony Calanducci
INFN Catania

CYCLOPS First Training Workshop
Bologna, 11h-13th April 2007

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

Metadata concepts and terminology

- Entries – Entities/Objects which we are attaching metadata to
- Attribute – key/value pair
  - 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
- Metadata - List of attributes (including their values) associated with entries
- Collection – A set of entries associated with a schema
  - *(AMGA collections are hierarchical organized)*
- Analogy to the RDBMS world: think of **collections** as **tables**, **attributes** as **columns**, **entries** as **rows**

Example: Movie Trailers

- Movie trailers files (entries) saved on Grid Storage Elements and registered into a LFC File Catalogue
- We have a LFN (Logical File Name) per movie files
- We want to add metadata to describe movie content. Possible schema:
  - Title -- varchar
  - Runtime -- int
  - Cast -- varchar
  - LFN -- varchar
- AMGA will be the repository of the movies' metadata

Screenshoots

```
Query> selectattr /trailers:Title Runtime FILE 'Runtime > 80'
>> Amelie of Montmartre
>> 122
>> 004405ac-da9a-1417-92db-c1ced08dbeef
>> American Pie 2
>> 108
>> 006d56b4-d7d1-1417-8417-c1ced08dbeef
>> Batman Begins
>> 141
>> 0072f510-db33-1417-b12e-c1ced08dbeef
>> The Fast and The Furious
>> 106
>> 00737e72-d8cb-1417-871f-c1ced08dbeef
>> Madagascar
>> 86
>> 0069b608-d95c-1417-9fd1-c1ced08dbeef
>> The Matrix
```

```
Query> ls
>> 004405ac-da9a-1417-92db-c1ced08dbeef
>> 006d56b4-d7d1-1417-8417-c1ced08dbeef
>> 0072f510-db33-1417-b12e-c1ced08dbeef
>> 00737e72-d8cb-1417-871f-c1ced08dbeef
>> 0069b608-d95c-1417-9fd1-c1ced08dbeef
>> 0010bf6c-d9cc-1417-a38c-c1ced08dbeef
>> 002e3966-d877-1417-8b9c-c1ced08dbeef
```

```
Query> listattr /trailers
>> Title
>> varchar(200)
>> Runtime
>> int
>> Country
>> varchar(25)
>> ReleaseDate
>> int
>> Director
>> varchar(80)
>> PlotOutline
>> text
>> Cast
>> varchar(2048)
>> Genre
>> varchar(100)
>> Image
>> text
```

gMOD: grid Movie On Demand

- gMOD provides a Video-On-Demand service
- User chooses among a list of video and the chosen one is streamed in real time to the video client of the user's workstation
- For each movie a lot of details (Title, Runtime, Country, Release Date, Genre, Director, Case, Plot Outline) are stored and users can search a particular movie querying on one or more attributes
- Two kind of users can interact with gMOD: TrailersManagers that can administer the db of movies (uploading new ones and attaching metadata to them); GILDA VO users (guest) can browse, search and choose a movie to be streamed.

gMOD under the hood

- Built on top of gLite services:
  - **Storage Elements**, sited in different place, physically contain the movie files
  - **LFC**, the File Catalogue, keeps track in which Storage Element a particular movie is located
  - **AMGA** is the repository of the detailed information for each movie, and makes possible queries on them
  - The **Virtual Organization Membership Service (VOMS)** is used to assign the right role to the different users
  - The **Workload Management System (WMS)** is responsible to retrieve the chosen movie from the right Storage Element and stream it over the network down to the user's desktop or laptop

gMOD interactions

VOMS

get Role

User

GENIUS Portal

Metadata Catalogue

AMGA

Storage Elements

LFC File Catalogue

Workload Management System

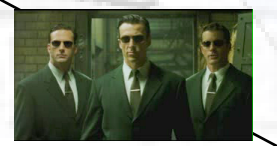

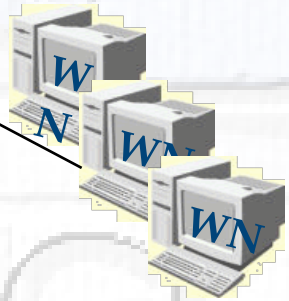

Information Society and Media

gMOD screenshot

gMOD is accesible through the GENIUS Portal (https://glite-tutor.ct.infn.it)

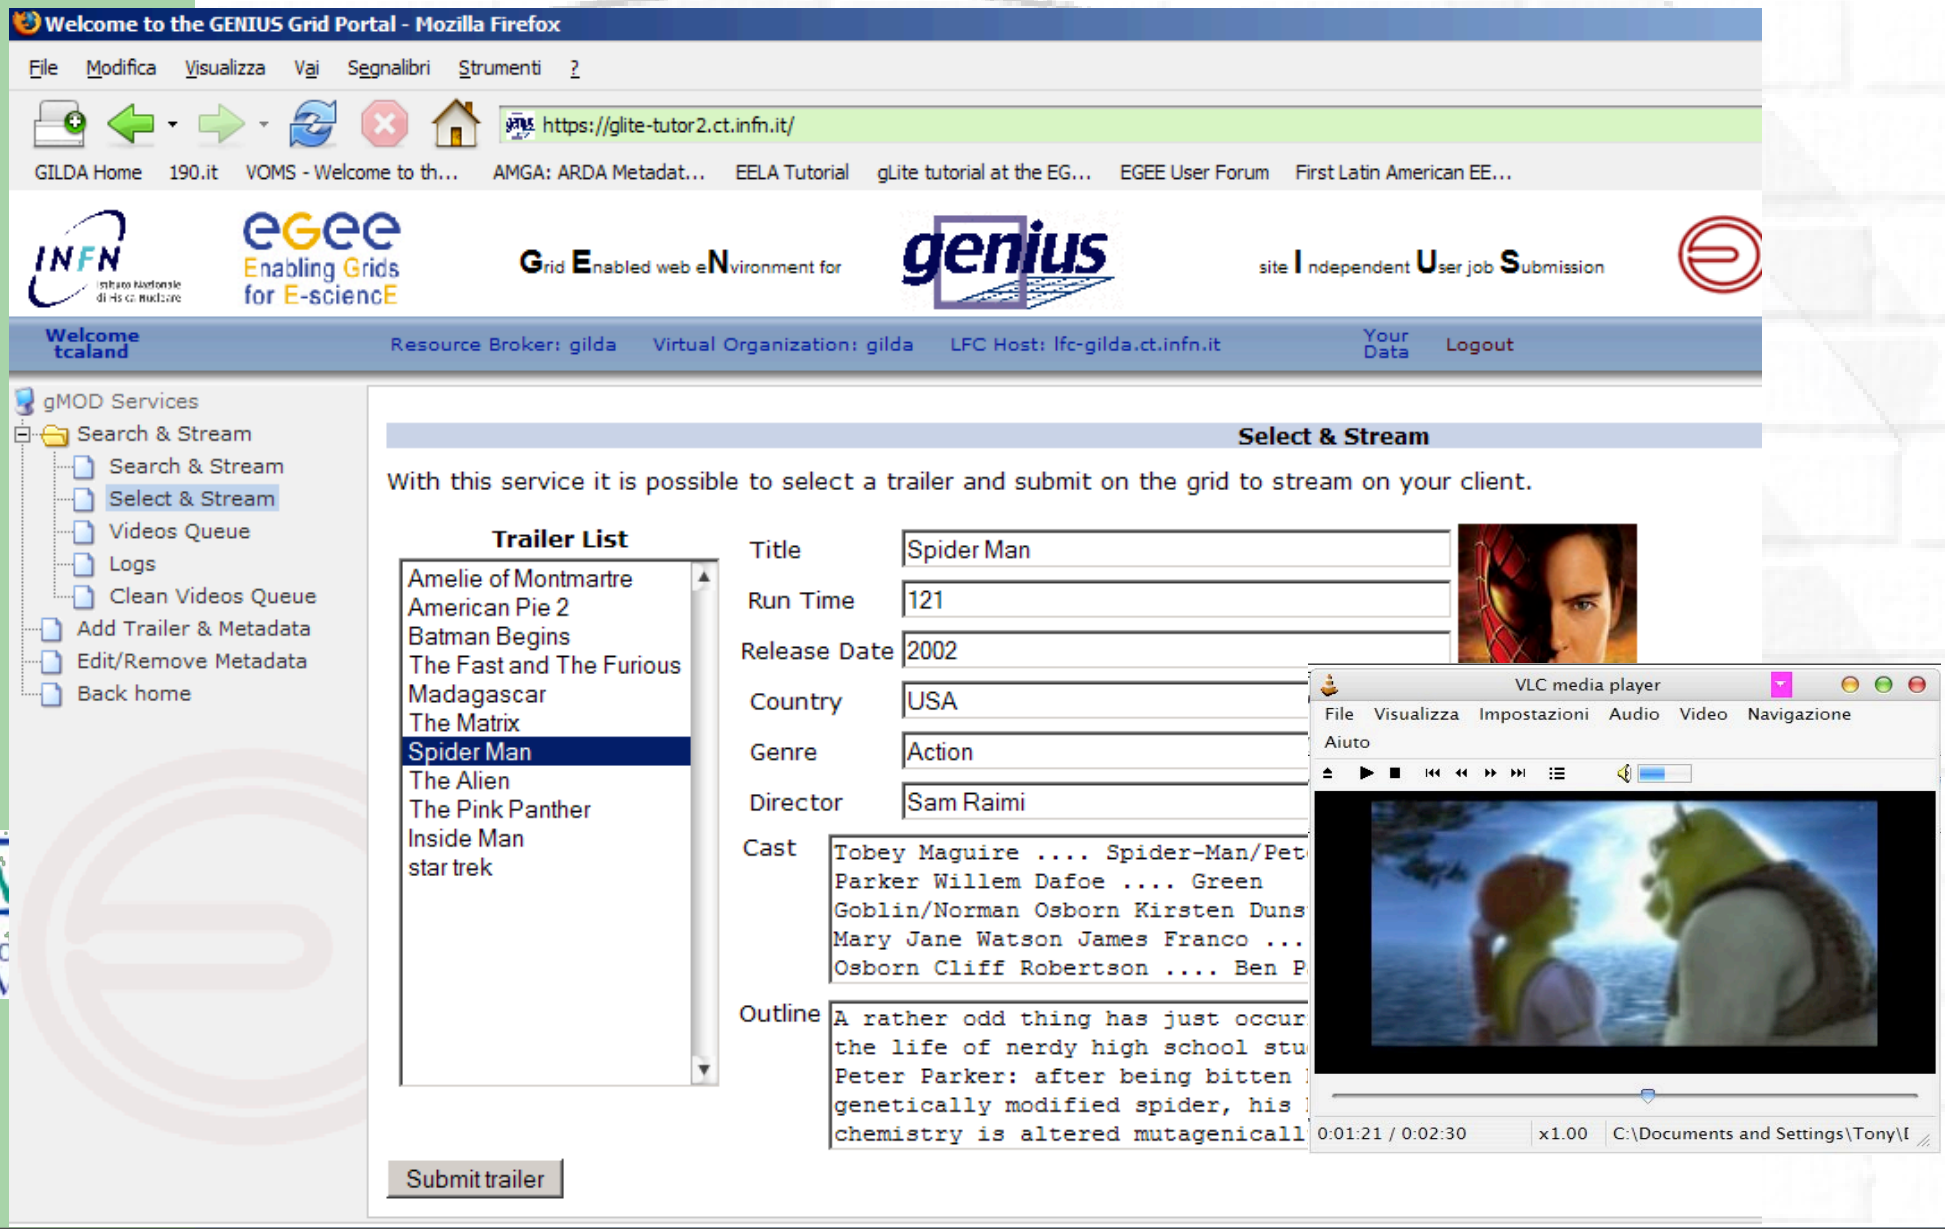

AMGA Features

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

```
selectattr /gLibrary:FileName /gLAudio:Author /gLAudio:Album
'/gLibrary:FILE=/gLAudio:FILE and like(/gLibrary:FileName, "%.mp3")'
```

AMGA 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):

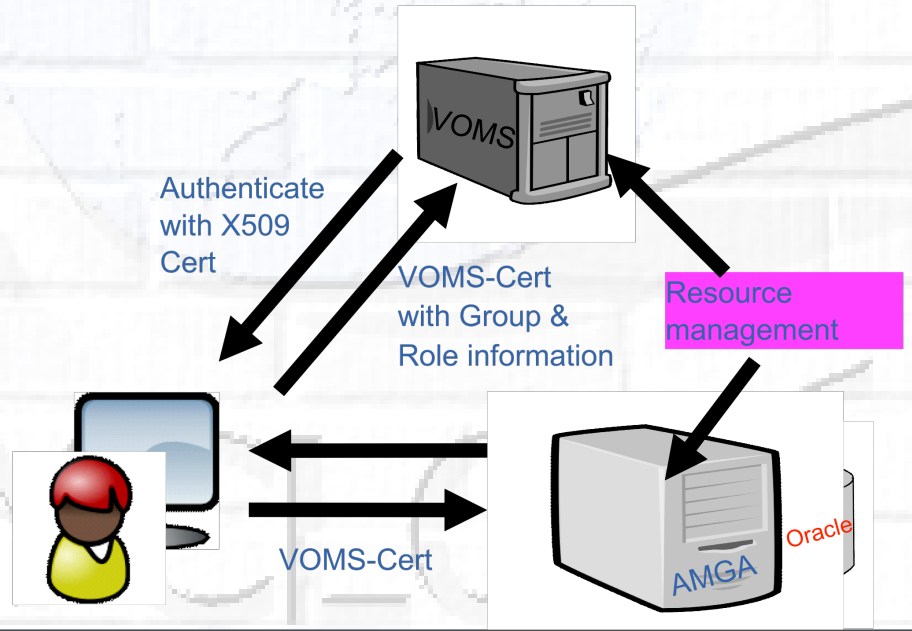

AMGA Implementation

- C++ multiprocess server
  - Runs on any Linux flavour
- Backends
  - Oracle, MySQL, PostgreSQL, SQLite
- Two frontends
  - TCP Streaming
    - High performance
    - Client API for C++, Java, Python, Perl, Ruby
  - SOAP
    - Interoperability
- Also implemented as standalone Python library
  - Data stored on filesystem

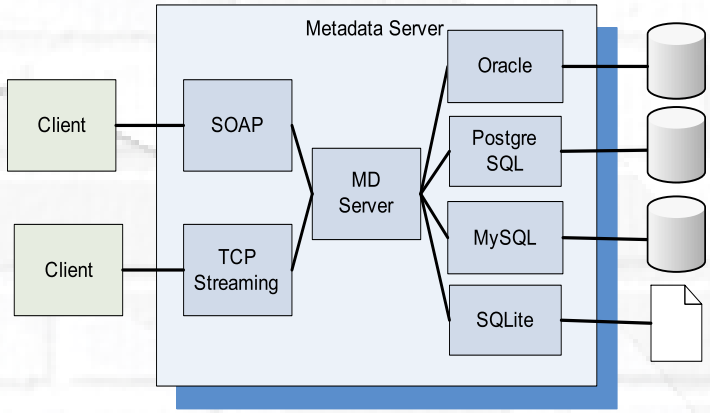

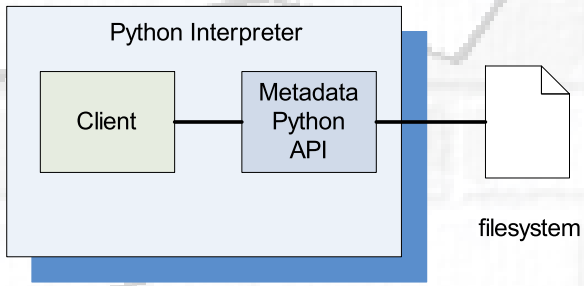

AMGA metadata types

AMGA Datatypes

	PostgreSQL	MySQL	Oracle	SQLite	Python
int	integer	int	number(38)	int	int
float	double precision	double precision	float	float	float
varchar(n)	character varying(n)	character varying(n)	varchar2(n)	varchar(n)	string
timestamp	timestamp w/o TZ	datetime	timestamp(6)	unsupported	time (unsupp.)
text	text	text	long	text	string
numeric(p,s)	numeric(p,s)	numeric(p,s)	numeric(p,s)	numeric(p,s)	float

- Using the above datatypes you are sure that your metadata can be easily moved to all supported back-ends
- If you do not care about DB portability, you can use, in principle, as entry attribute type ALL the datatypes supported by the back-end, even the more esoteric ones (PostgreSQL Network Address type or Geometric ones)

- We played a little bit with GIS Datatype offered by MySQL 5

Example with ESR data

We created a /ESR/opera_nno collection asking AMGA to use the MyISAM table engine

```
Query> listattr /ESR/opera_nno
>> Dataset
>> varchar(30)
>> File_Name
>> varchar(50)
>> Footprint
>> multipolygon
>> Lat
>> numeric(8,2)
>> Level
>> varchar(5)
```

```
>> Lon
>> numeric(8,2)
>> Orbit
>> int(5)
>> Proc_centre
>> varchar(50)
>> Proc_date
>> timestamp
>> Start_Date
>> timestamp
>> Stop_Date
>> timestamp
...
```

We used *insert* command that evaluates all inserted values:

```
insert sameEntryName Dataset "GOME" Level 2 Version "v1.1" Orbit 25421 File_Name "/
grid/esr/gome/utv/2000/03/00301000.utv" Start_Date '"2000-02-29 00:01:00.0"' Stop_Date
'"2000-02-29 00:58:00.0"' Footprint 'MPolyFromText("MULTIPOLYGON(((82.96 -59.12,75.95
-89.07,75.95 -89.07,76.46 -94.77,76.84 -100.85,77.07 -107.21,77.13 -115.34,77.00
-121.80,76.72 -128.08,76.30 -134.03,75.74 -139.59,75.07 -144.70,74.30 -149.36,80.26
-179.07,80.26 -179.07,81.52 -174.78,82.71 -169.12,83.81 -161.42,84.76 -150.74,85.47
-136.17,85.80 -117.93,85.57 -94.31,84.94 -78.84,84.03 -67.39,82.96 -59.12)))")'
Proc_centre "EGEE" Proc_date '"2005-10-14 13:20:00.0"' File_input "00301000.lv1"
Proc_description '"Algorithm: utv"'
```

Sample queries

Let's check if the entry was properly inserted (we need to use AsText() to decode a MultiPolygon):

```
Query> selectattr /ESR/opera_nno:File_Name AsText(/ESR/opera_nno:Footprint) ''
>> /grid/esr/gome/utv/2000/03/00301000.utv
>> MULTIPOLYGON(((82.96 -59.12,75.95 -89.07,75.95 -89.07,76.46 -94.77,76.84
-100.85,77.07 -107.21,77.13 -115.34,77 -121.8,76.72 -128.08,76.3 -134.03,75.74
-139.59,75.07 -144.7,74.3 -149.36,80.26 -179.07,80.26 -179.07,81.52 -174.78,82.71
-169.12,83.81 -161.42,84.76 -150.74,85.47 -136.17,85.8 -117.93,85.57 -94.31,84.94
-78.84,84.03 -67.39,82.96 -59.12)))
```

We want to look for a Polygon that cointains a given point:

```
Query> selectattr /ESR/opera_nno:File_Name /ESR/opera_nno:Start_Date /ESR/
opera_nno:Stop_Date 'Contains(/ESR/opera_nno:Footprint, GeomFromText("POINT(82.96
-59.12)"))'
>> /grid/esr/gome/utv/2000/03/00301000.utv
>> 2000-02-29 00:01:00
>> 2000-02-29 00:58:00
```

- As a summary, the following functions work: GeomFromText(), MPolyFromText(), Contains(), AsText()
- In principle PostgreSQL+PostGIS would also work but this is not tested.

Accessing AMGA from UI/WNs

- TCP Streaming Front-end
  - mdcli & mdclient and C++ API (md_cli.h, MD_Client.h)
  - Java Client API and command line mdjavaclient.sh & mdjavacli.sh (also under Windows !!)
  - Python and Perl Client API
  - **PHP Client API – NEW**
    - **eveloped totally by the GILDA team – INFN CT**
  - **AMGA Web Interface** ---NEW
    - Developed totally by the GILDA team – INFN CT
    - Based on JAVA AMGA Standard APIs
    - Web Application using standard as JSP Custom Tags, Servlet
- SOAP Frontend (WSDL)
  - C++ gSOAP
  - AXIS (Java)
  - ZSI (Python)

AMGA Web Interface

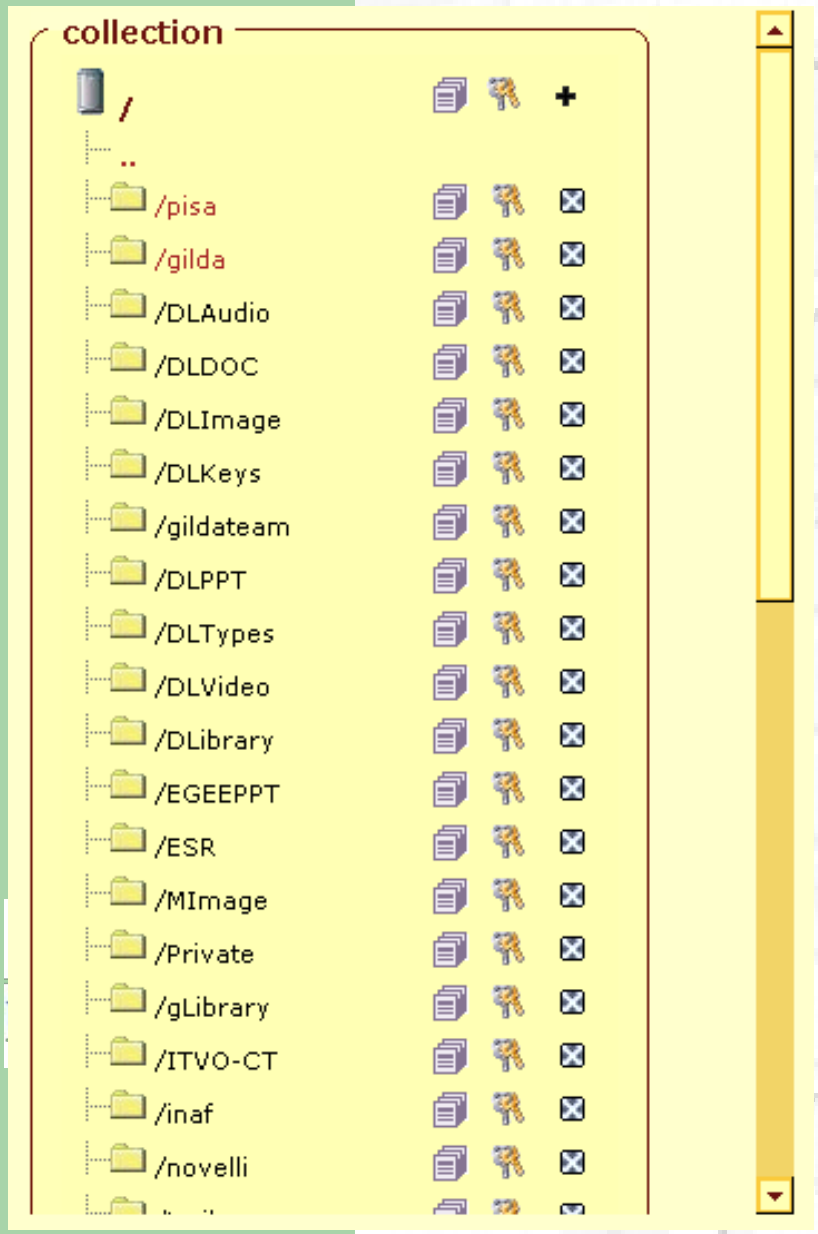

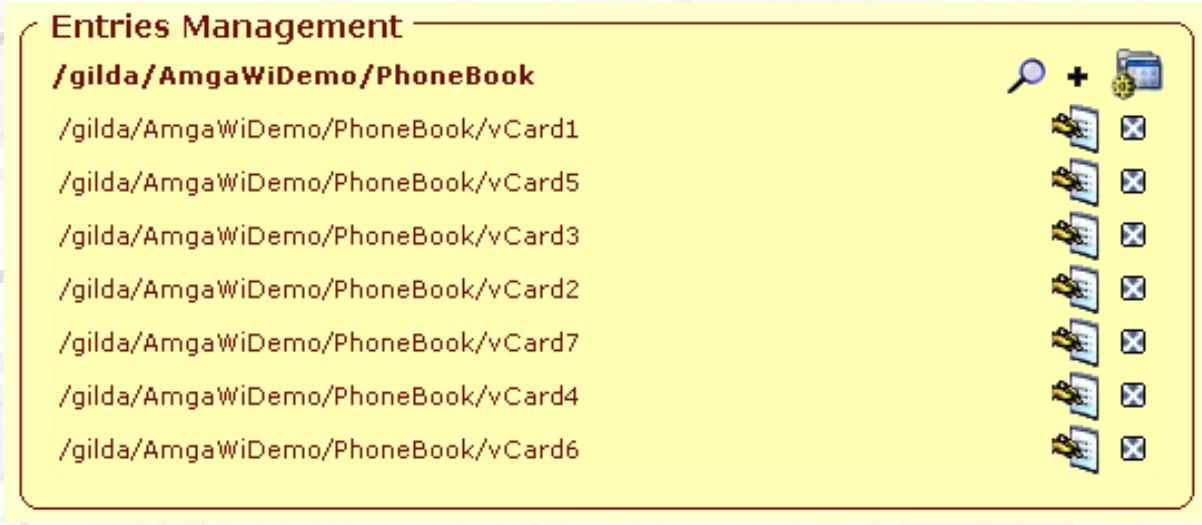

Metadata Schema Management

Metadata Replication

- 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

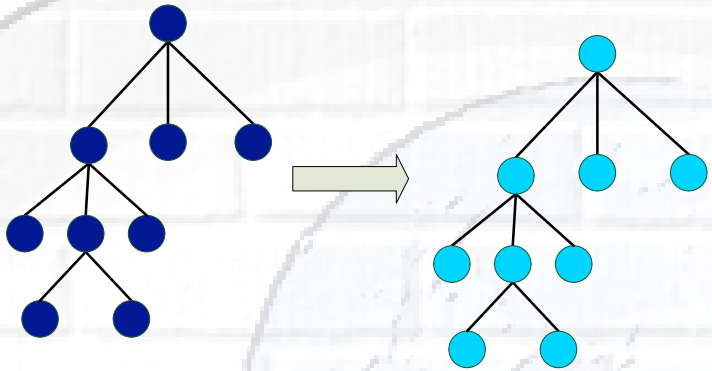

Partial replication

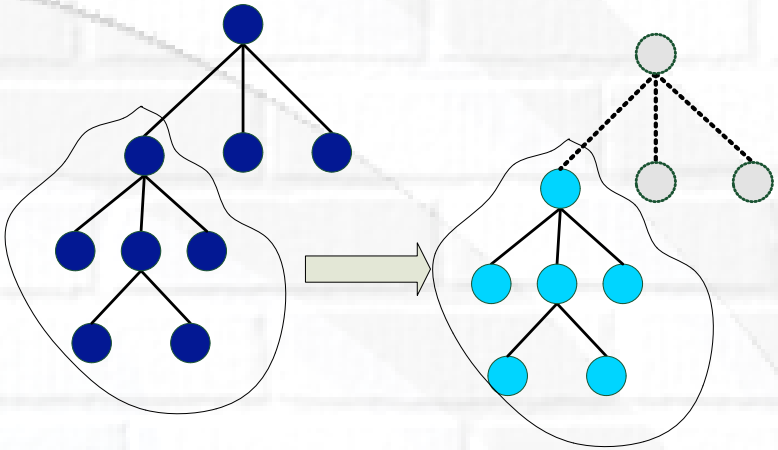

Federation

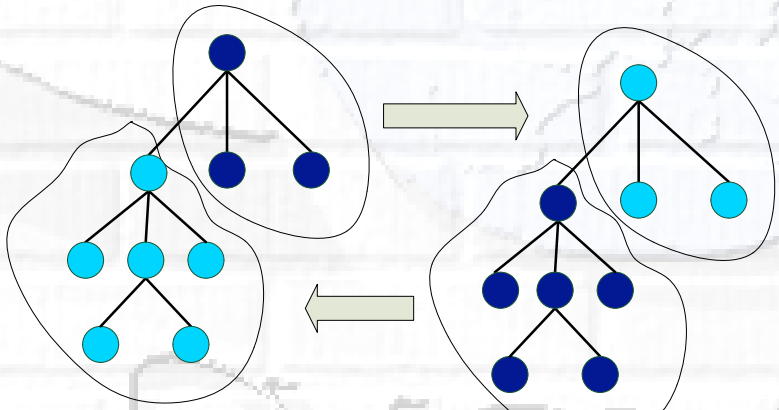

Proxy

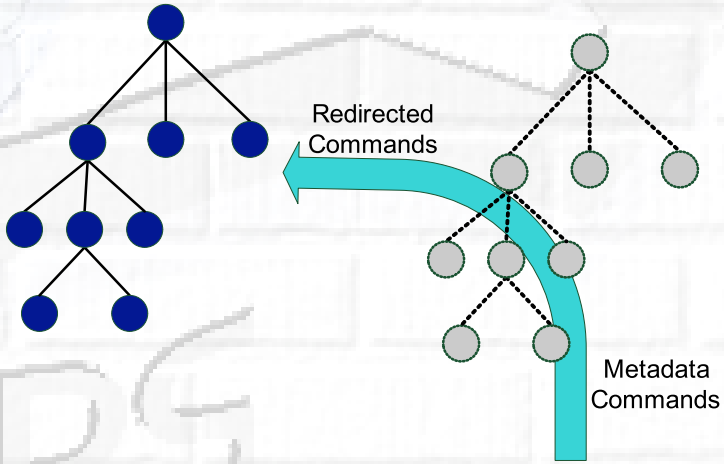

Early adopters of AMGA

- 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

Use AMGA to exchange data among running jobs

- Suppose we have two sets of jobs:
  - Producers: they generate a file, store on a SE, register it onto the LFC File Catalogue assigning a LFN
  - Consumers: they will take a LFN, download the file and elaborate it

- AMGA can be used to share the information generated by the Producers, it could act as a "bag-of-LFNs" (bag-of-task model) from which Consumers can fetch file for further elaboration

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

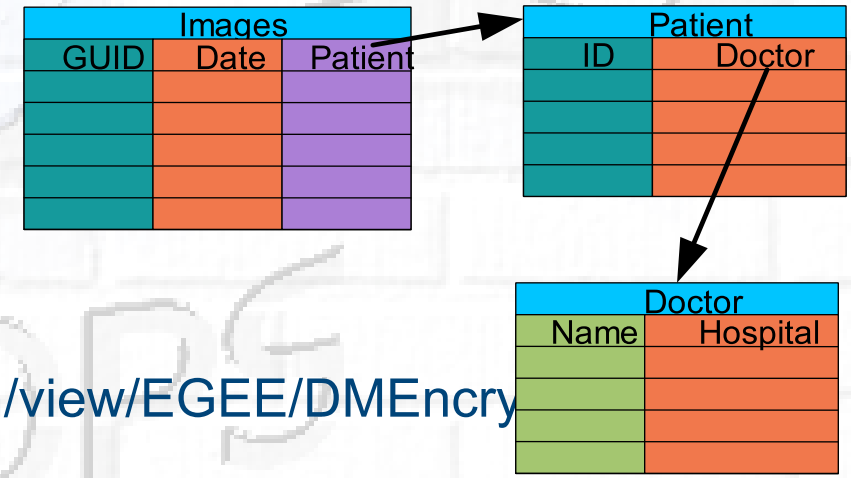

- More details at
  - https://uimon.cern.ch/twiki/bin/view/EGEE/DMEncry

gLibrary

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?

CYCLOPS

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?
  - (if you have good memory, its GUID could be a solution ☺)

Information Society and Media

CYCLOPS

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?
  - (if you have good memory, its GUID could be a solution ☺)
  - File Catalogues just let us to arrange files in folders and subfolders, no way to *query on their contents*

Information Society and Media

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?
  - (if you have good memory, its GUID could be a solution ☺)
  - File Catalogues just let us to arrange files in folders and subfolders, no way to *query on their contents*
  - Metadata Catalogues are a possible solution, but not always “affordable” especially for non expert users (powerful but complex to use)

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?
  - (if you have good memory, its GUID could be a solution ☺)
  - File Catalogues just let us to arrange files in folders and subfolders, no way to *query on their contents*
  - Metadata Catalogues are a possible solution, but not always “affordable” especially for non expert users (powerful but complex to use)
- Our solution: a higher level application built on top of many gLite grid services: a Metadata Catalogue + File Catalogues + Storage Elements →

gLibrary

- Huge amounts of data can be saved on SEs (did we forget about the existence of Data Grids?)
- But how can we easily find later a file that we need?
  - (if you have good memory, its GUID could be a solution ☺)
  - File Catalogues just let us to arrange files in folders and subfolders, no way to *query on their contents*
  - Metadata Catalogues are a possible solution, but not always “affordable” especially for non expert users (powerful but complex to use)
- Our solution: a higher level application built on top of many gLite grid services: a Metadata Catalogue + File Catalogues + Storage Elements →
- Requirements: easy to use, fast, secure, extensible

gLibrary goals

gLibrary goals

- Attempt to create a Digital Asset Management System for the Grid
  - Examples of Digital Assets handled by gLibrary:
    - Images
    - Videos
    - Audio Files
    - Office Documents (Powerpoint, Word, Excel, OpenOffice)
    - E-Mails, PDFs, HTMLs
    - Customized versions of the previous well-know document type (ex. EGEE PPTs)
    - ….
- Keep track and organize in a uniform way all the additional details (metadata) of files saved in Storage Elements and registered in File Catalogues
- Provide users an easy way to locate and retrieve files based on their contents

CYCLOPS

Some usage scenarios

Information Society and Media

Some usage scenarios

- Example 1:
  - Locate all theoretical (PPTType) PowerPoint (Type) presentations about “gLite DMS” (Keywords) given in 2005 (Date) by Uncle Sam (Speaker);
  - Find all the movies (Type) in which Julia Roberts (Cast) performed together with Hugh Grant (Cast) produced in USA (Country) in 2004 (ReleaseDate); or all the acoustic (Genre) mp3 (Format) audio files (Type) of Alanis Morissette (Singer) that last more than 3 minutes (Runtime).

CYCLOPS

Some usage scenarios

- Example 1:
  - Locate all theoretical (PPTType) PowerPoint (Type) presentations about “gLite DMS” (Keywords) given in 2005 (Date) by Uncle Sam (Speaker);
  - Find all the movies (Type) in which Julia Roberts (Cast) performed together with Hugh Grant (Cast) produced in USA (Country) in 2004 (ReleaseDate); or all the acoustic (Genre) mp3 (Format) audio files (Type) of Alanis Morissette (Singer) that last more than 3 minutes (Runtime).
- Example 2:
  - A doctor is looking for brain (keyword) DICOM (Type) images of male (Gender) patients older than 65 (Age).

Information Society and Media

CYCLOPS

Some usage scenarios

- Example 1:
  - Locate all theoretical (PPTType) PowerPoint (Type) presentations about “gLite DMS” (Keywords) given in 2005 (Date) by Uncle Sam (Speaker);
  - Find all the movies (Type) in which Julia Roberts (Cast) performed together with Hugh Grant (Cast) produced in USA (Country) in 2004 (ReleaseDate); or all the acoustic (Genre) mp3 (Format) audio files (Type) of Alanis Morissette (Singer) that last more than 3 minutes (Runtime).
- Example 2:
  - A doctor is looking for brain (keyword) DICOM (Type) images of male (Gender) patients older than 65 (Age).
- Example 3:
  - A job can behave as a storage crawler: it scans pre-existing files in Storage Elements to extract relevant metadata that will be published on gLibrary for further data mining.

Information Society and Media

Some gLibrary features

- Hierarchical types. Ex:
  - Audio
    - Music
    - Ringtones
    - SoundEffects
  - Video
    - Movies
    - Trailers
    - Clips
- Intuitive browsing (a la iTunes) with 3 customizable filter fields. Ex:
  - for Music you can browse by Genre, Artist, Album, Year, Rating, Format)
  - for Movie you can browse by Genre, ReleaseDate, Studio, Country, Director)
- Grouping of assets by Categories
  - to put together assets of different types but belonging to the same category (think for ex to all the files needed in a given project: images, ppt, pdf, sounds)
  - to narrow furtherly the assets of a given type (ex.: music playlists, preferred movies)
- Your digital libraries accessible from everywhere through a Web 2.0 frontend (AJAX based + Java Applets + PHP 5)

gLibrary Security

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)

Summary of Metadata Services

- AMGA – Metadata Service of gLite
  - Part of gLite (but still not certificed in gLite 3.0. it will be done with 3.1 release)
  - Useful for simplified DB access
  - Integrated on the Grid environment (Security)
- Replication/Federation features
- Tests show good performance/scalability
- Already deployed by several Grid Applications
  - LHCb, ATLAS, Biomed, …
  - AMGA WI, ADAT, gMOD, gLibrary

References

- AMGA Project Homepage
  - http://amga.web.cern.ch/amga/
- AMGA 1.2.7 User Manual
  - http://amga.web.cern.ch/amga/downloads/amga-manual_1_2_7.pdf
- Exercise documentation from ISSGC'06:
  - http://www.dma.unina.it/~murli/ISSGC06/glite/public_html/amga.html
  - http://www.dma.unina.it/~murli/ISSGC06/glite/public_html/summaryExercise.htm
- AMGA GILDA Wiki pages:
  - https://grid.ct.infn.it/twiki/bin/view/GILDA/AMGAHandsOn
  - https://grid.ct.infn.it/twiki/bin/view/GILDA/AMGAAdv

Questions…