Metadata Management with AMGA

Antonio Calanducci
antonio.calanducci@ct.infn.it
National Institute of Nuclear Physics
INFN Catania

Workshop CCR e INFN Grid

Santa Tecla (Acireale), 17-21 Maggio 2010

https://grid.infn.it/

Outline

- Why Metadata on the Grid
- AMGA Features
- Use cases
- gLibrary: digital libraries on the Grid

Why Grid needs Metadata?

- **Grids allow to save millions of files spread over several storage sites.**
- **Users and applications need an efficient mechanism**
  - to describe files
  - to locate files based on their contents
- **This is achieved by**
  - associating descriptive attributes to files
    - Metadata is data about data
  - answering user queries against the associated information

Trailer example

Entry names	Title	Ru	Cast	LFN
8c3315c1-811f-4823-a778-60a203439689	My Best Friend's	80	Julia Roberts	lfn:/grid/gilda/movies/mybfwed.avi
51a18b7a-fd21-4b2c-aa74-4c53ee64846a	Spider-man 2	120	Kirsten Dunst	lfn:/grid/gilda/movies/spiderman2.avi
401e6df4-c1be-4822-958c-ce3eb5c54fcb	The God Father	113	Al pacino	lfn:/grid/gilda/movies/godfather.avi

Trailer example

Attribute

Entry names	Title	Ru	Cast	LFN
8c3315c1-811f-4823-a778-60a203439689	My Best Friend's	80	Julia Roberts	lfn:/grid/gilda/movies/mybfwed.avi
51a18b7a-fd21-4b2c-aa74-4c53ee64846a	Spider-man 2	120	Kirsten Dunst	lfn:/grid/gilda/movies/spiderman2.avi
401e6df4-c1be-4822-958c-ce3eb5c54fcb	The God Father	113	Al pacino	lfn:/grid/gilda/movies/godfather.avi

Trailer example

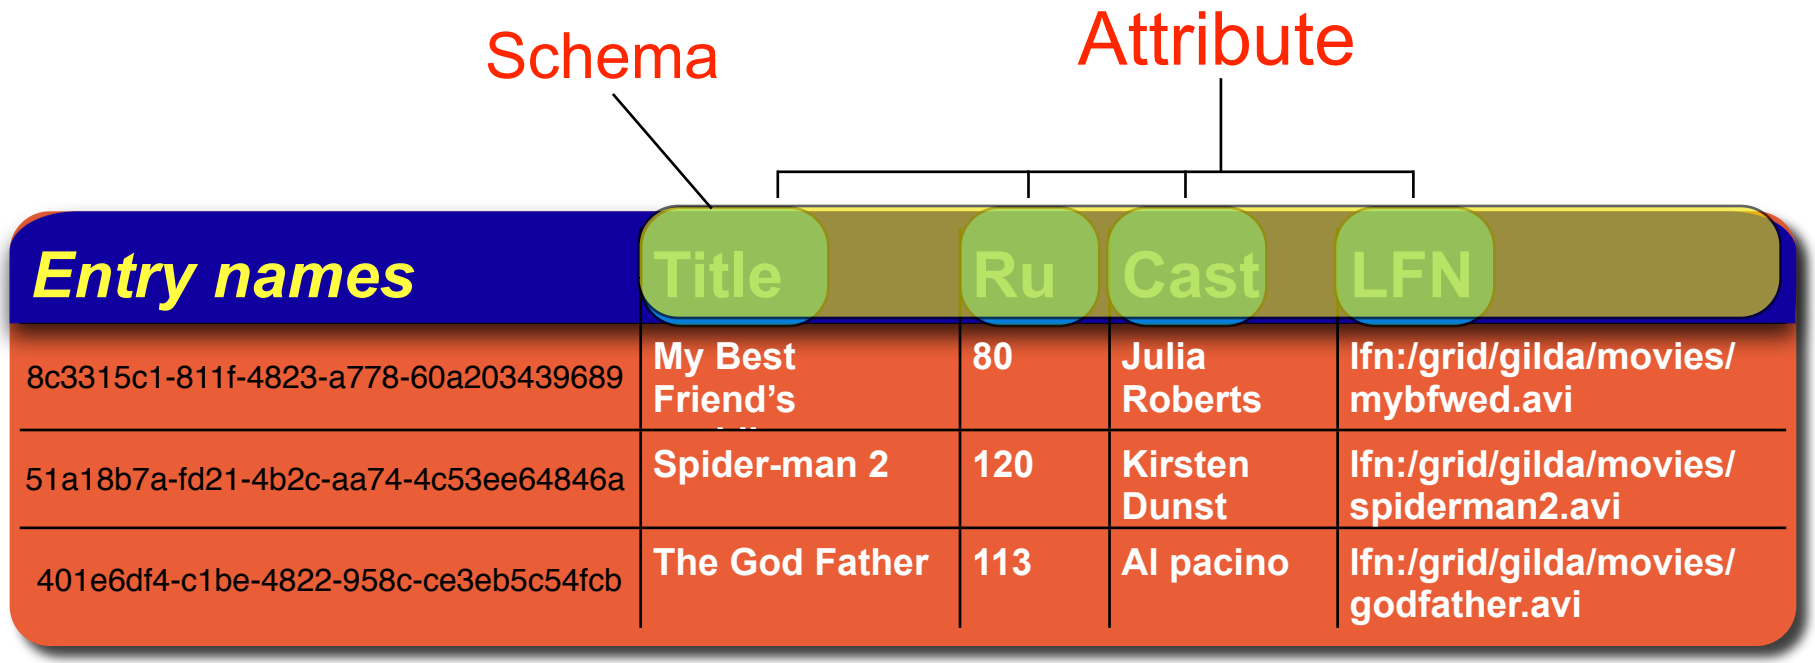

Trailer example

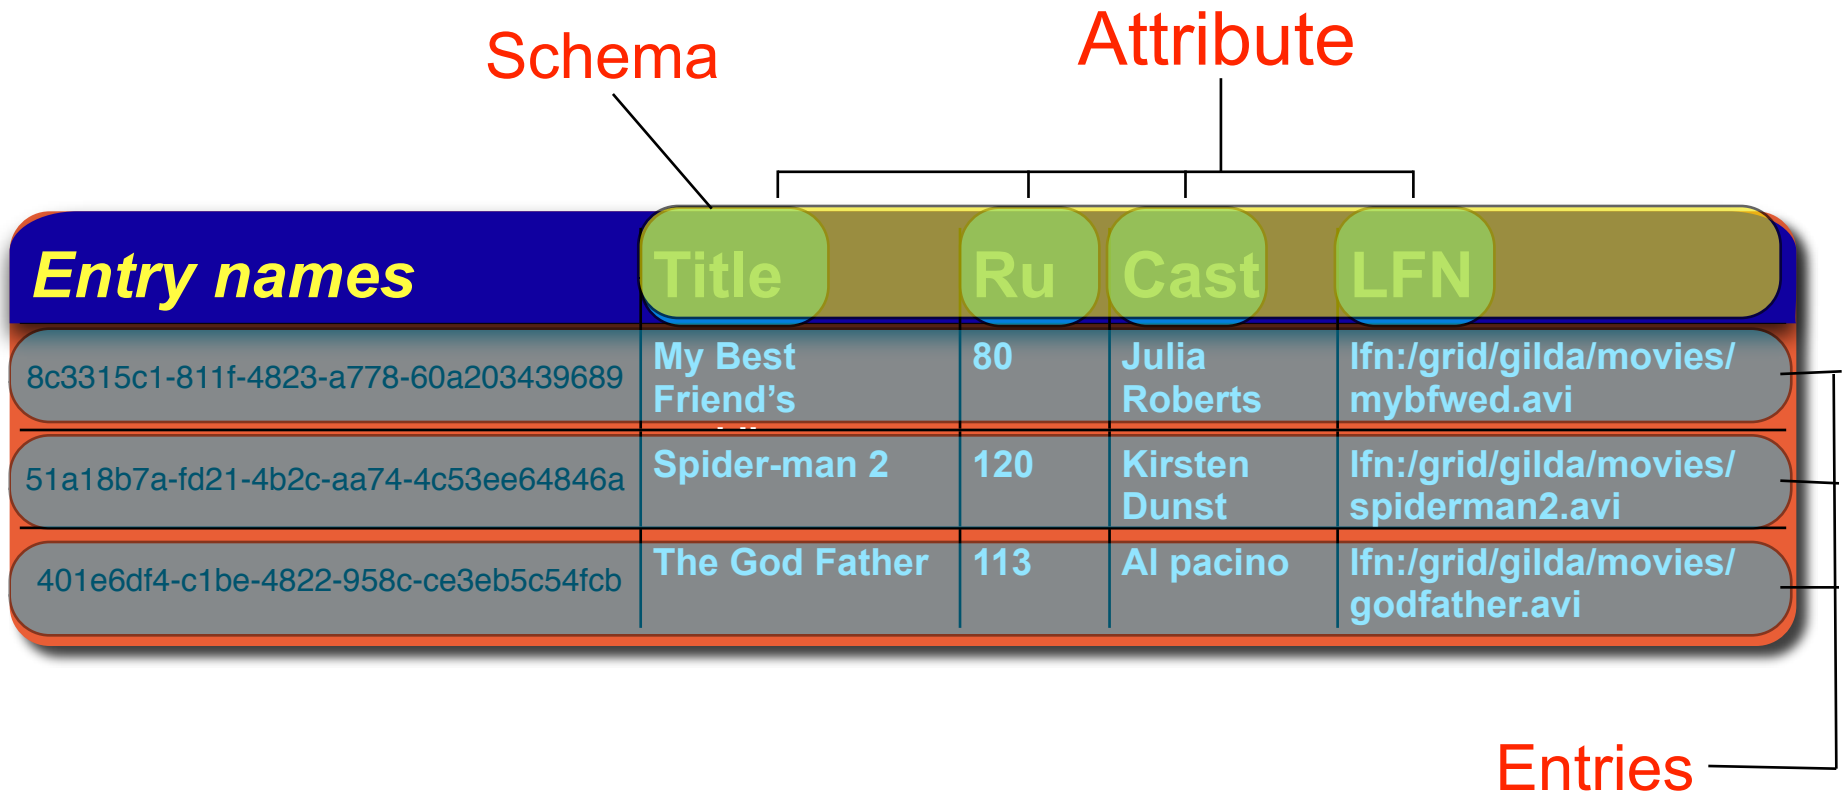

Trailer example

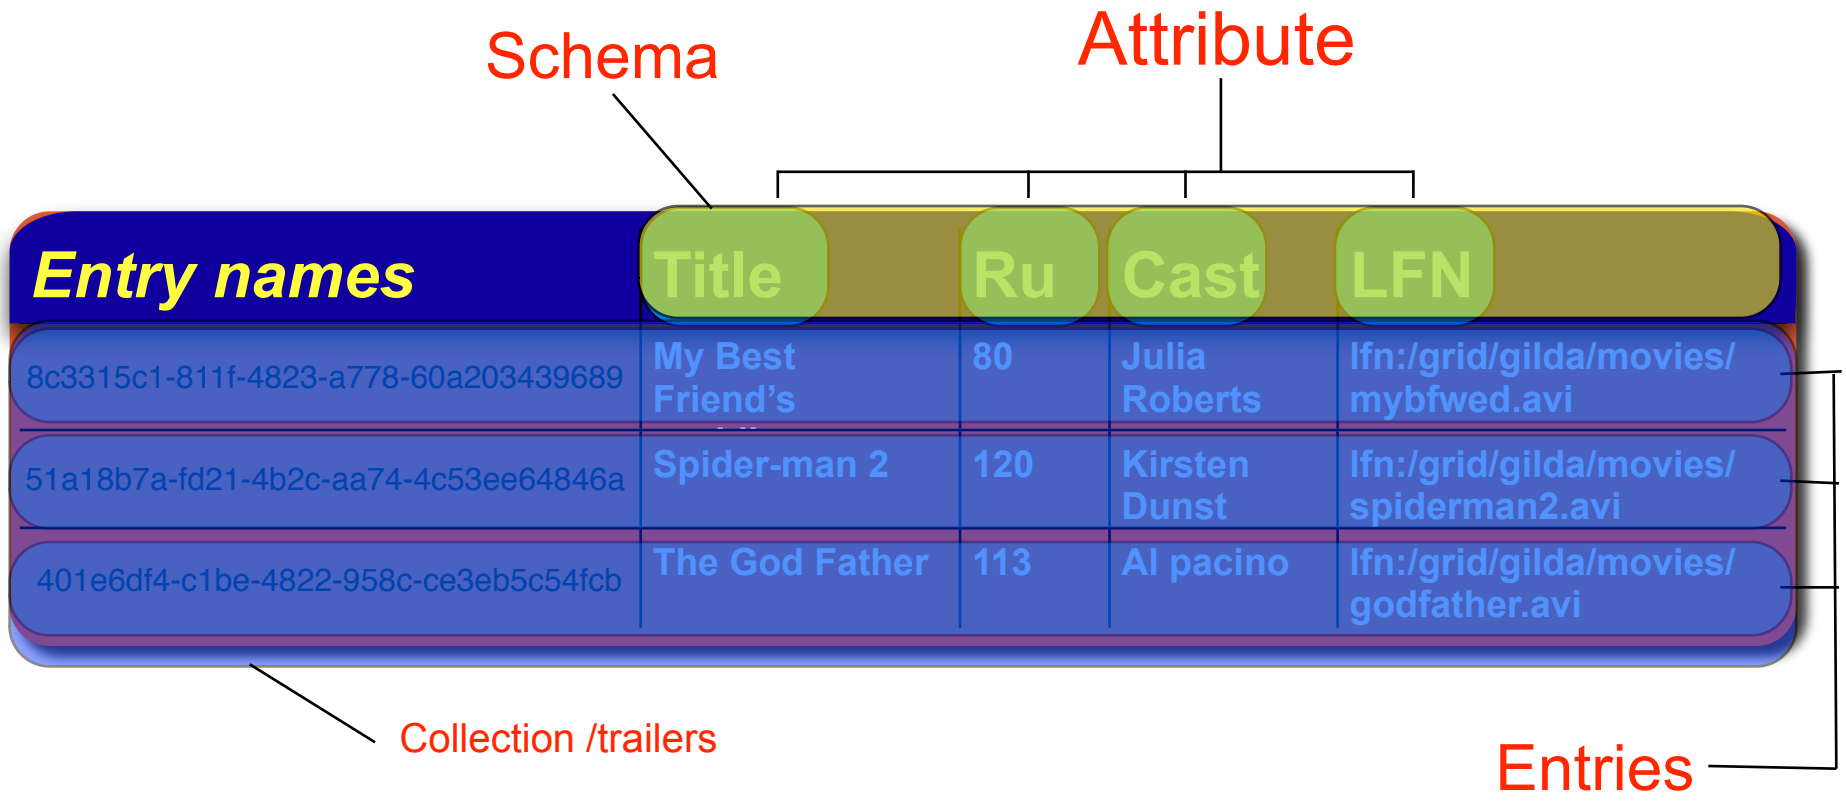

Metadata service on the Grid

Metadata service on the Grid

- **Information about files -- but not only!**

Metadata service on the Grid

- **Information about files -- but not only!**
- **metadata can describe any grid entity/object**
  - ex: JobIDs - add logging information to your jobs

Metadata service on the Grid

- **Information about files -- but not only!**
- **metadata can describe any grid entity/object**
  - ex: JobIDs - add logging information to your jobs
- **monitoring of running applications:**
  - ex: ongoing results from running jobs can be published on the metadata server

Metadata service on the Grid

- **Information about files -- but not only!**
- **metadata can describe any grid entity/object**
  - ex: JobIDs - add logging information to your jobs
- **monitoring of running applications:**
  - ex: ongoing results from running jobs can be published on the metadata server
- **Inputset for a storm of parametric jobs**

Metadata service on the Grid

- **Information about files -- but not only!**
- **metadata can describe any grid entity/object**
  - ex: JobIDs - add logging information to your jobs
- **monitoring of running applications:**
  - ex: ongoing results from running jobs can be published on the metadata server
- **Inputset for a storm of parametric jobs**
- **information exchanging among grid peers**
  - ex: producers/consumers job collections: master jobs produce data to be analyzed; slave jobs query the metadata server to retrieve input to “consume”

Metadata service on the Grid

- **Information about files -- but not only!**
- **metadata can describe any grid entity/object**
  - ex: JobIDs - add logging information to your jobs
- **monitoring of running applications:**
  - ex: ongoing results from running jobs can be published on the metadata server
- **Inputset for a storm of parametric jobs**
- **information exchanging among grid peers**
  - ex: producers/consumers job collections: master jobs produce data to be analyzed; slave jobs query the metadata server to retrieve input to “consume”
- **Simplified DB access on the grid**
  - Grid applications that needs structured data can model their data schemas as metadata

Monitoring of running application

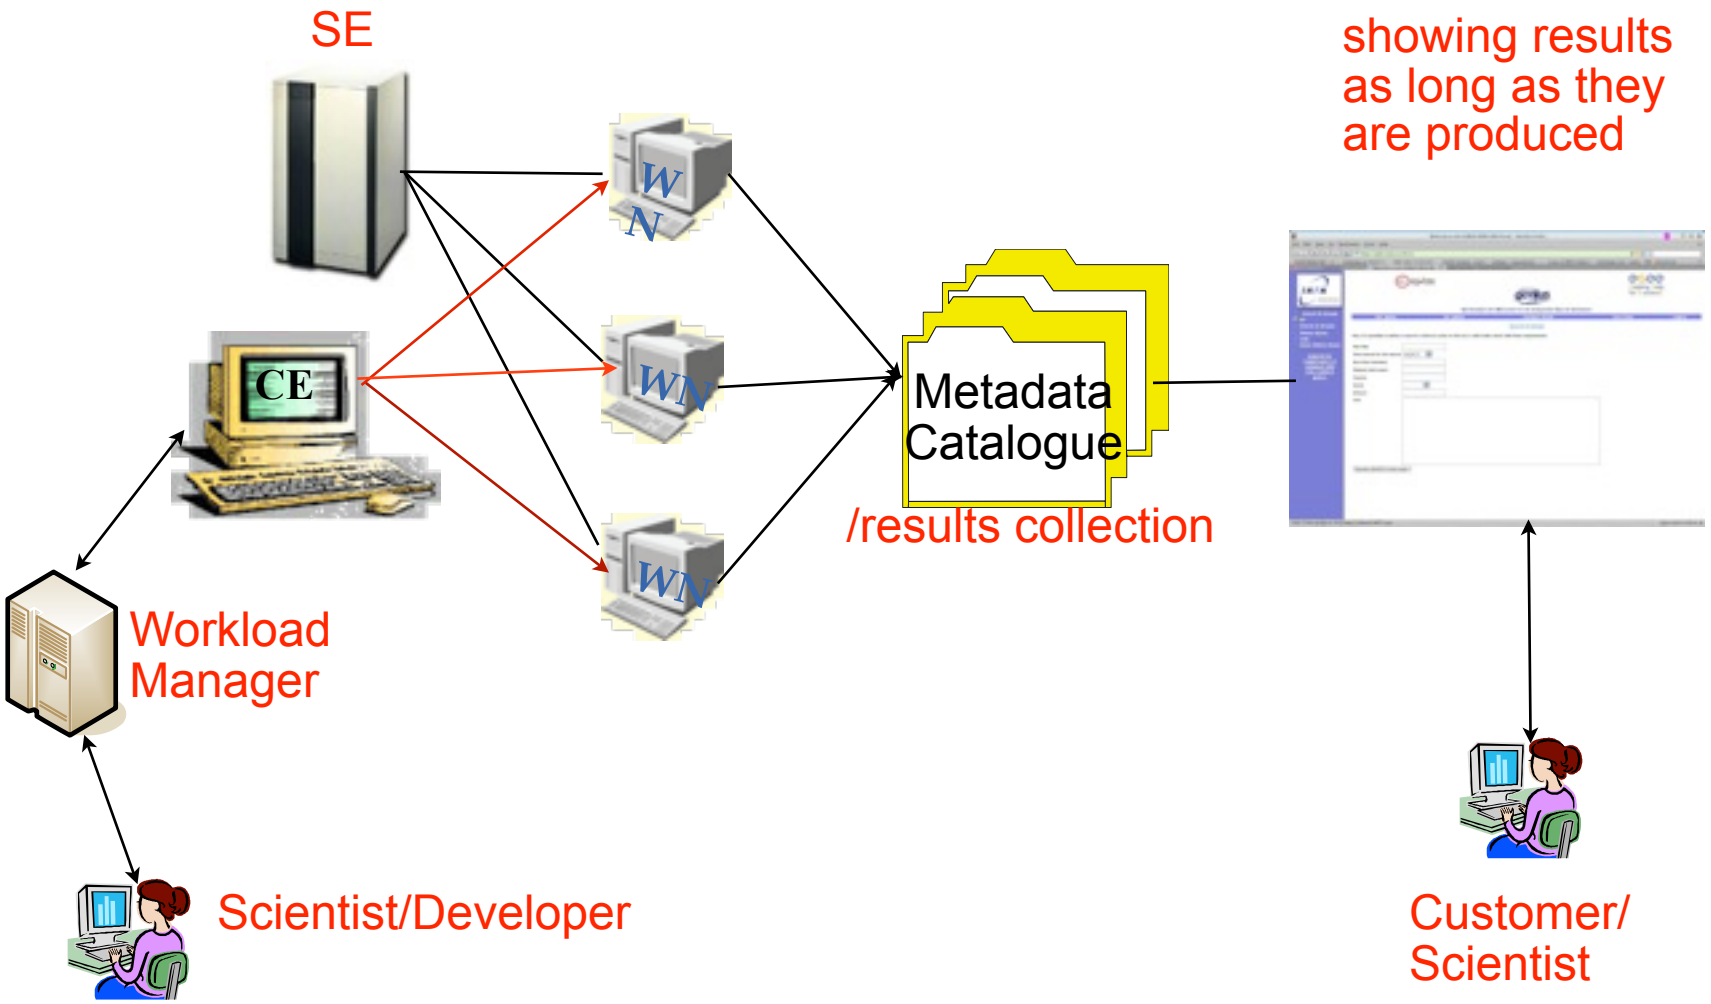

Inputset for parametric jobs

- /grid/my_simulation/input

entry	x1	x2	y1	y2	step	isTaken	found	output
1	9453.1	9453.32	-439.93	-439.91	0.0006	JobID1234	No pillars	
2	9342.13	3435	3423	2343.2	0.003	No		
3	34254.3	342342	432.43	132	0.002	No		
..... and so on								

- This collection lists all the parameter set to be run on the Grid
- On the WN, one of the inputset is selected and “isTaken” is set = JOB_ID of the job that has fetched it
- Results is also written in the “found” column to monitor the simulation
  - so users can check the simulation from a UI, querying the metadata server, or from a WebPage (using APIs for ex)
- StdOutput can be copied also into the “output” text column

Information exchanging among grid peers

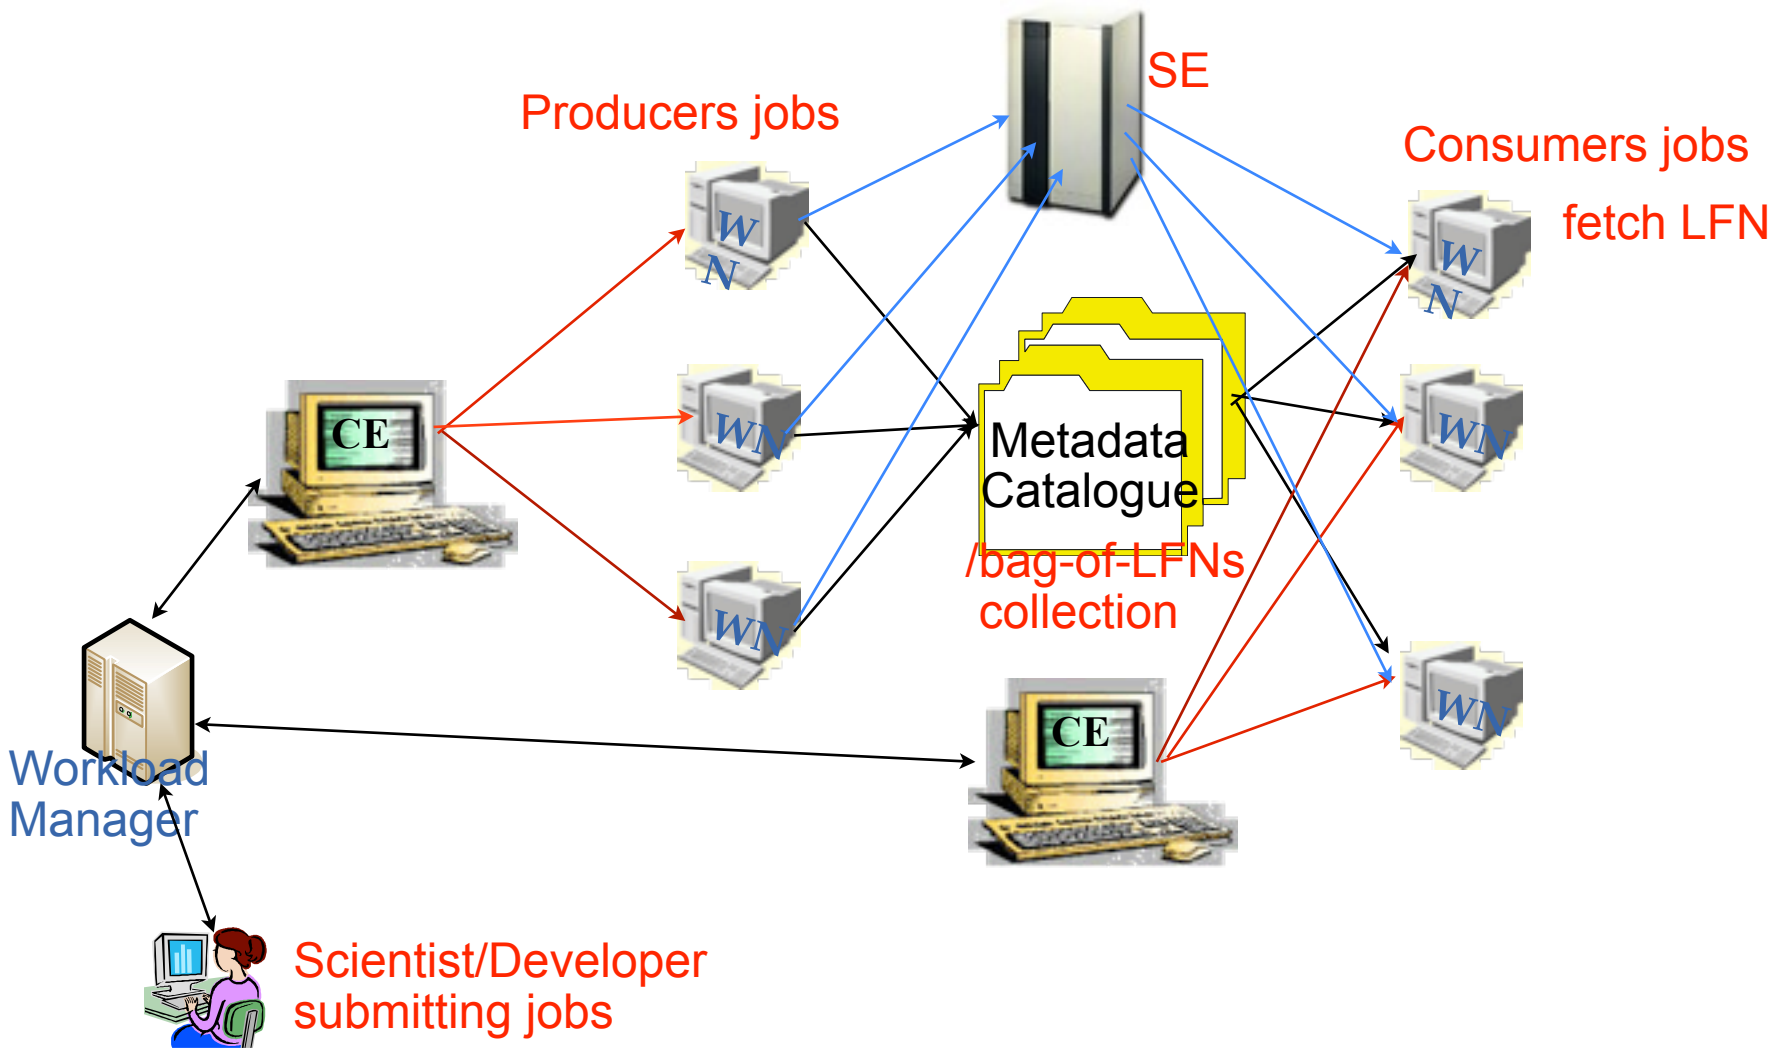

The AMGA Metadata Catalogue

- **Official metadata service for the gLite middleware**
  - but no dependencies from gLite software
  - **it can be used with other grid technologies/other environments**
- **AMGA: Arda Metadata Grid Application**
- **Provide a complete but simple interface, in order to make all users able to use it easily.**
- **Designed with scalability in mind in order to deal with large number of entries**
  - based on a lightweight and streamed text-based protocol, like HTTP/SMTP
- **Grid security is provided to grant different access levels to different users.**
- **Flexible with support to dynamic schemas in order to serve several application domains**
- **Simple installation by tar source, RPMs or Yum/YAIM**

AMGA Security

AMGA Security

- Unix style permissions - users and groups

AMGA Security

- Unix style permissions - users and groups
- ACLs – Per-collection or per-entry (table row).

AMGA Security

- Unix style permissions - users and groups
- ACLs – Per-collection or per-entry (table row).
- Secure client/server connections – SSL

AMGA Security

- **Unix style permissions - users and groups**
- **ACLs – Per-collection or per-entry (table row).**
- **Secure client/server connections – SSL**
- **Client Authentication based on**
  - Username/password
  - General X509 certificates (DN based)
  - Grid-proxy certificates (DN based)

AMGA Security

- **Unix style permissions - users and groups**
- **ACLs – Per-collection or per-entry (table row).**
- **Secure client/server connections – SSL**
- **Client Authentication based on**
  - Username/password
  - General X509 certificates (DN based)
  - Grid-proxy certificates (DN based)
- **VOMS support:**
  - VO attribute maps to defined AMGA user
  - VOMS Role maps to defined AMGA user
  - VOMS Group maps to defined AMGA group

AMGA Implementation

- **C++ multiprocess server**
  - **Backends**
    - Oracle, MySQL 4/5, PostgreSQL, SQLite
  - **Front Ends**
    - **TCP text streaming**
      - *High performance*
      - *Client API for C++, Java, Python, Perl, PHP*
    - **SOAP (deprecated)**
      - *Interoperability*
      - *Scalability*
    - **WS-DAIR Interface (new in AMGA 2.0)**
      - ***WS-enable environment***
- **Standalone Python Library implementation**
  - Data stored on file system

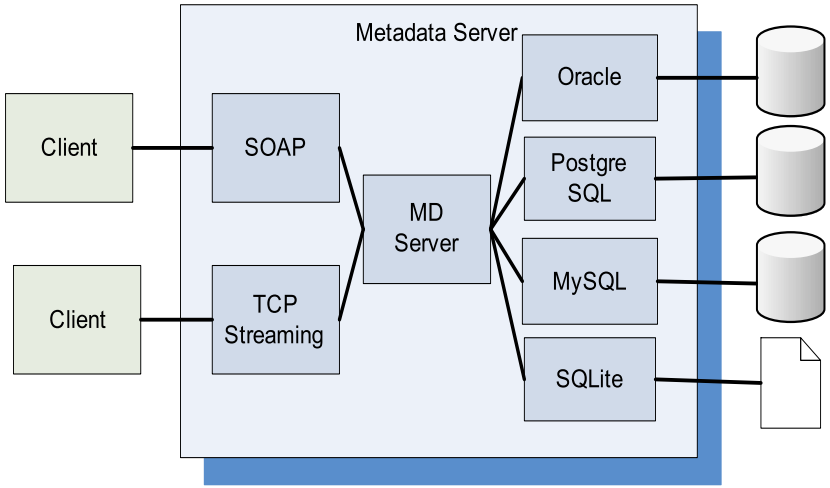

- AMGA server runs on SLC3/4, Fedora Core, Gentoo, Debian

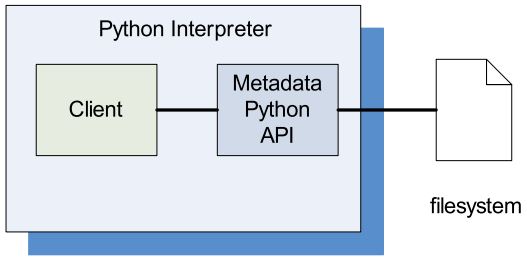

Metadata Replication: Use cases

Full replication

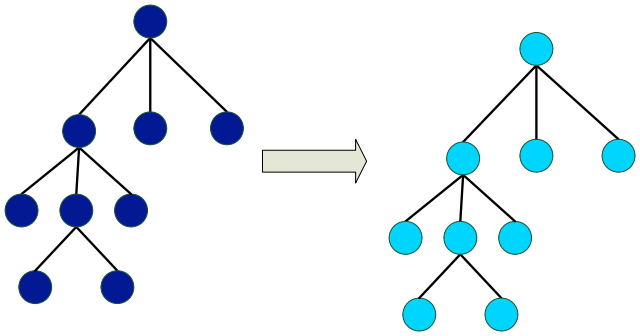

Partial replication

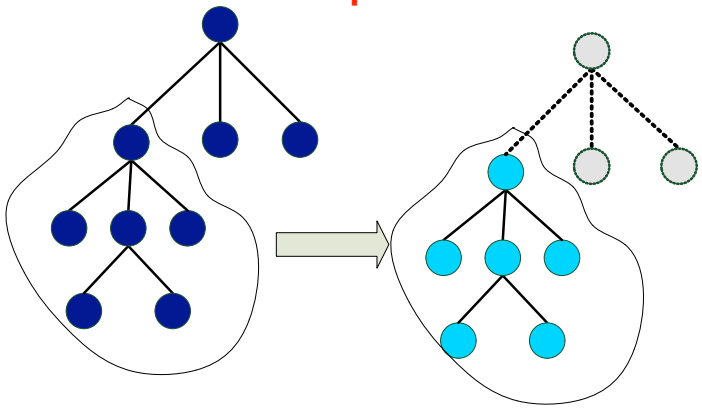

Federation

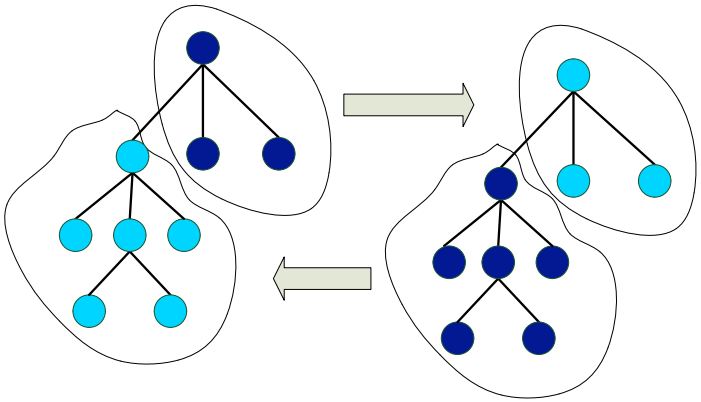

Proxy

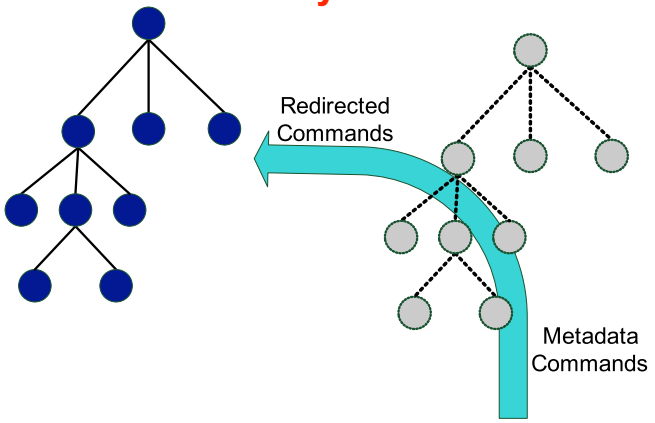

Existing SQL DBs importing

- **Since AMGA 1.2.10, a new import feature allow to access existing DB table**
- **Once imported into AMGA the tables from one or more DBs you want to access through AMGA, you can exploit many of the features brought to you by AMGA for your existing tables**
- **Advantages:**
  - your db tables can be accessed by grid users/applications, using grid authentication (VOMS proxies)/authorization with ACLs
  - exploiting AMGA federation features you can access several databases together from the Grid

Native SQL Support

- **Objective:**
  - implement native SQL query processing functionality in AMGA
- **Current Status:**
  - direct SQL data statement in SQL92 Entry Level has been implemented in the 1.9 release
    - Including 4 statements: SELECT, DELETE, UPDATE and INSERT
    - ALL SQL commands should be issued in UPPERCASE
- **Entry name:**
  - when a new entry is created with addentry/addentries, a name has to be assigned (filling the “file” column in the AMGA db backend)
    - in the INSERT implementation, it’s filled automatically with a random guid

Early adopters of AMGA

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

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

Biomed - MDM

- **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)

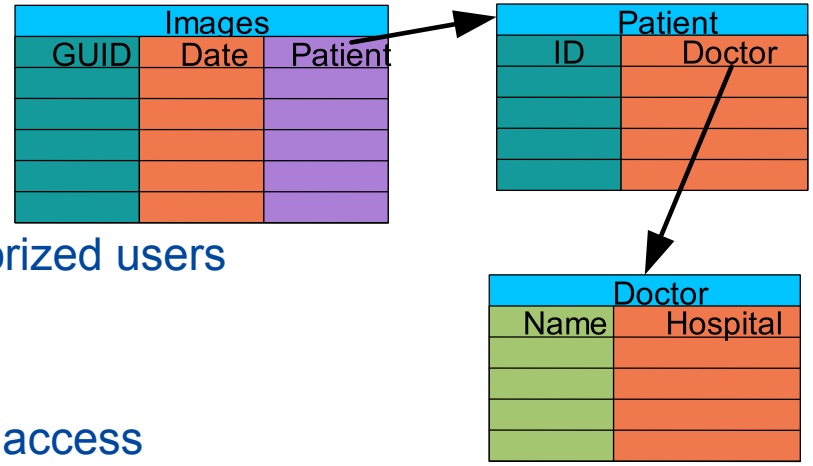

- **Strong security requirements**
  - Patient data is sensitive
  - Data must be encrypted
  - Metadata access must be restricted to authorized users
- **AMGA used as metadata server**
  - Demonstrates authentication and encrypted access
  - Used as a simplified DB
- **More details at**
  - https://uimon.cern.ch/twiki/bin/view/EGEE/DMEncryptedStorage

gLibrary features

- INFN-developed tool totally gLite based
- It allows to store, organize, search and retrieve digital assets on a Grid environment with an intuitive front-end
- What we mean by **Digital Assets**:

gLibrary as the iTunes for the Grid

Browse & Search

- **Assets can be browsed selecting a type (or category) and selecting one or more filters:**
  - attributes of the selected types, chosen from a defined list, used to narrow the result set
- **Filter application is cascading and context-sensitive: the selection of a filter value dynamically influences subsequent filter values (*"à la iTunes"* browsing)**
  - Classical search by description and keywords available too

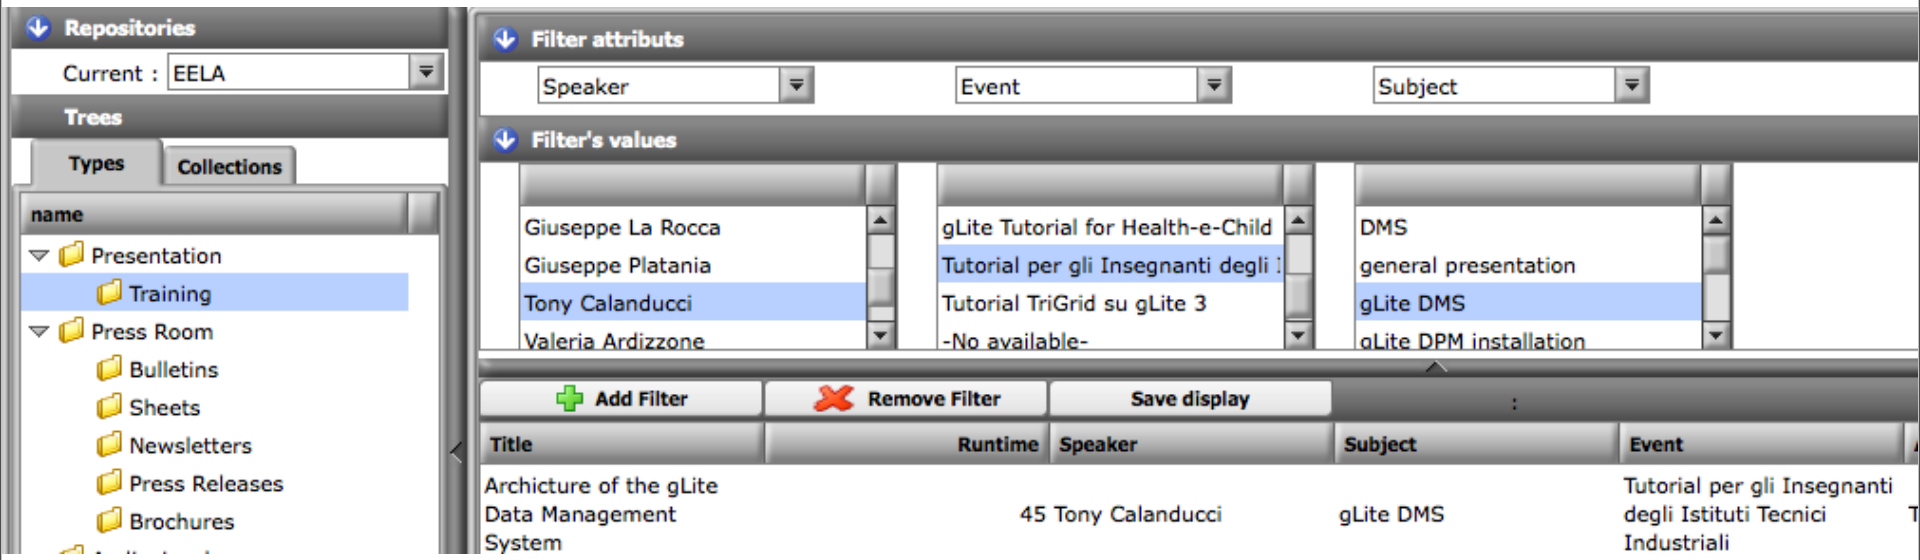

Organize assets

- **"Types" and "Categories" definition by repository providers/admins:**
- **Assets are organized by *type*:**
  - a list of specific attributes to describe each kind of asset to be managed by the system
  - hierarchical (a child type shares and extend parent's attributes)
  - queried during searches

EXAMPLE OF TYPES AND ATTRIBUTES' LIST

Type	Attributes' list
Audio	Format, Bitrate, Samplerate, Time
Music	(Format, Bitrate, Samplerate, Time), Name, Artist, Album, Genre, Tracknumber, Year, Artwork, Lyric, Rating
Presentation	Format, NumOfPages
Training	(Format, NumOfPages), Title, Runtime, Speaker, Author, Subject, Event, Date, Type
(Root)	*FileName, SubmissionDate, Description, Keywords, LastModificationDate, Size*

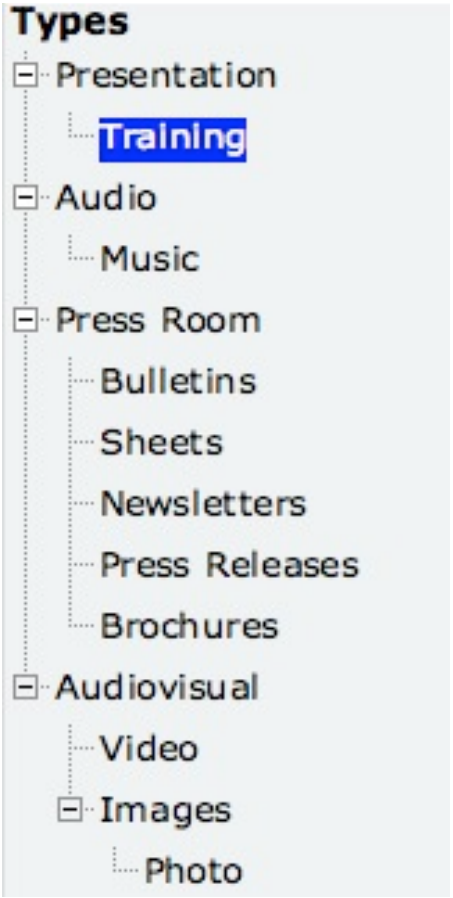

- **and/or organized by collection:**
  - Group together related assets of different type;
  - Useful also to define subsets of assets belonging to the same type
  - Multiple category assignment per asset (tagging)

Store & Retrieve

- **Users can upload their local assets on one or more (creating replicas) Storage Elements of the Grid**
  - Files already on grid SE can be registered in a gLibrary repository by the LFC File Catalogue browser

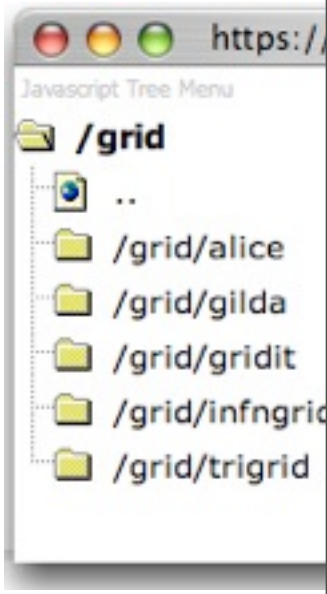

- **Download from SEs to the users' laptop/desktop:**
  - selection of a replica link from a list

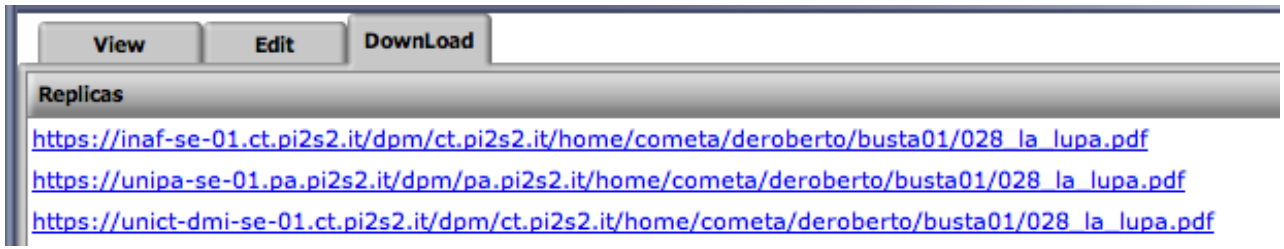

- **Transfers are handled from the browser over HTTP/HTTPS provided that users have their own X.509 Grid Certificate imported**

gLibrary architecture

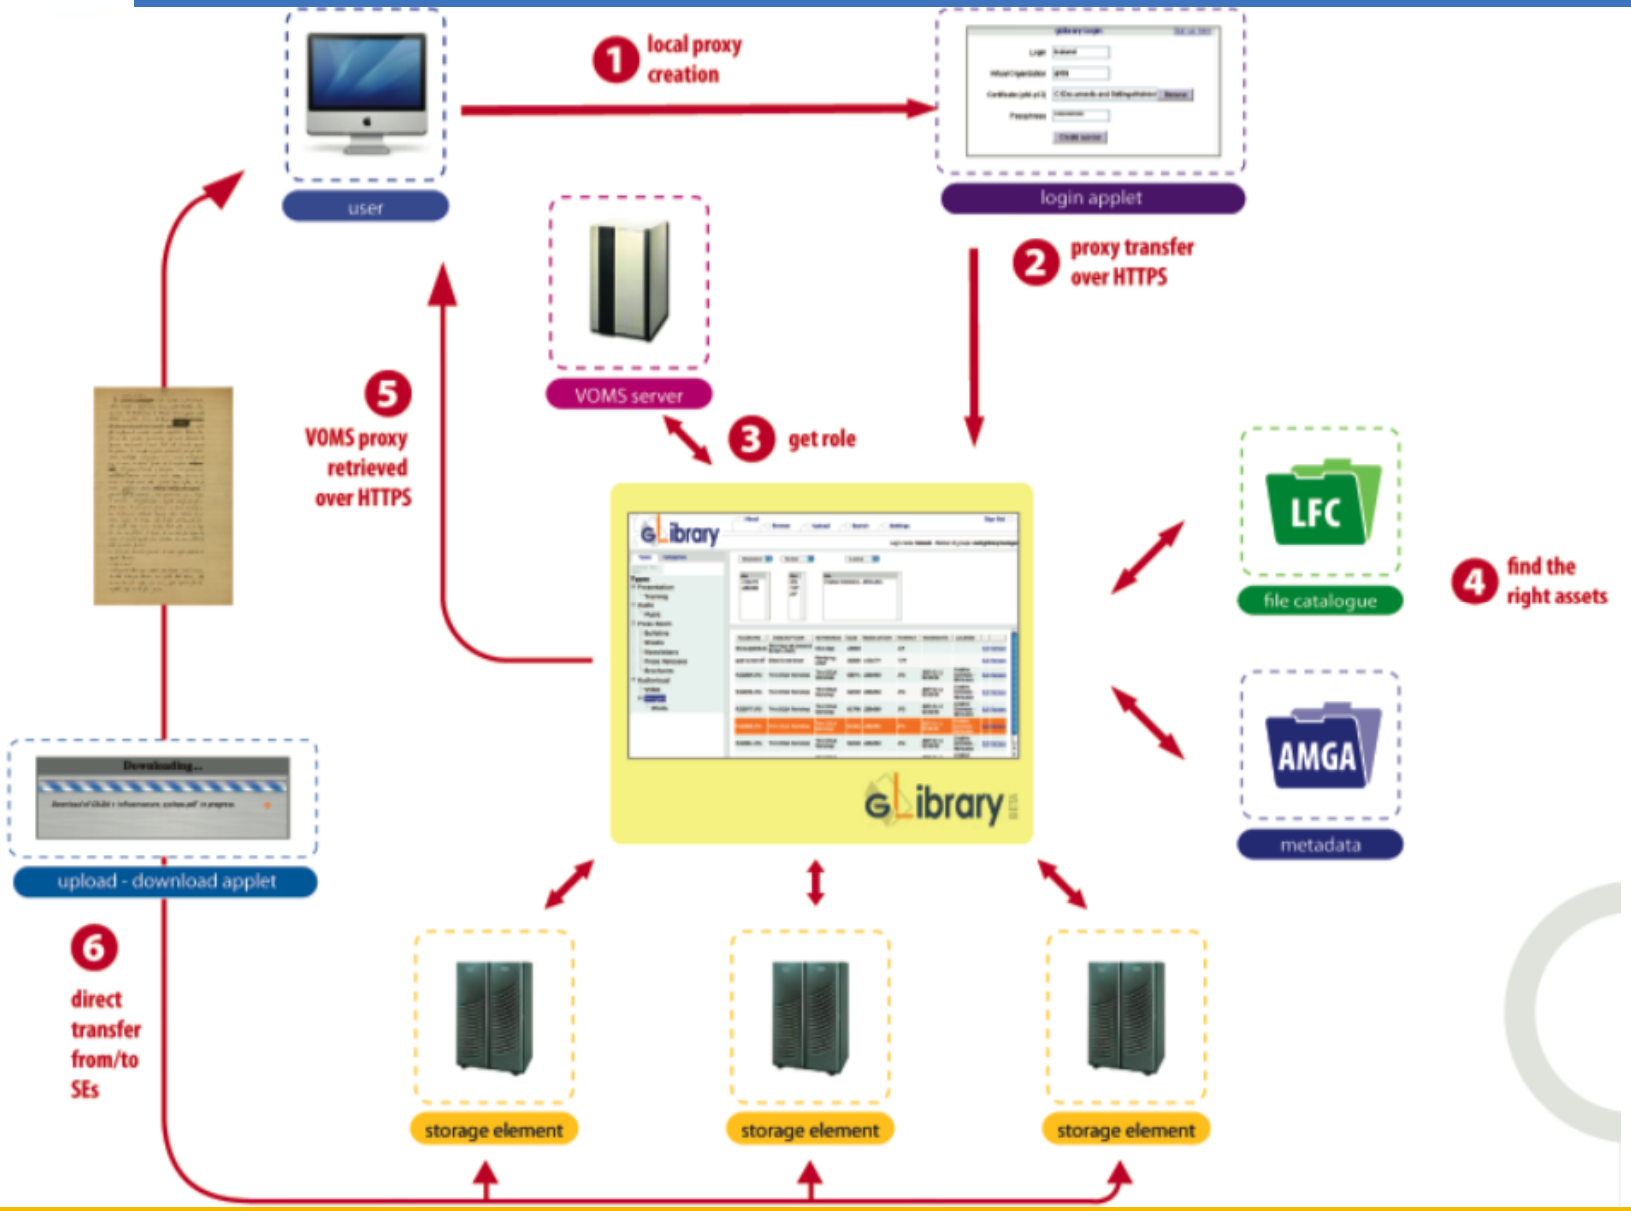

De Roberto cultural heritage

- De Roberto, an Italian writer of the XIX/XX century, born in Naples, but spending his life in Catania, has left to the humanistic community numerous works
- Those are made up of valuable and hard-to-manage pieces: manuscripts, typescripts, drafts with handwritten corrections, magazines, cuts, sketches, photos, etc.

Il repository di De Roberto con gLibrary

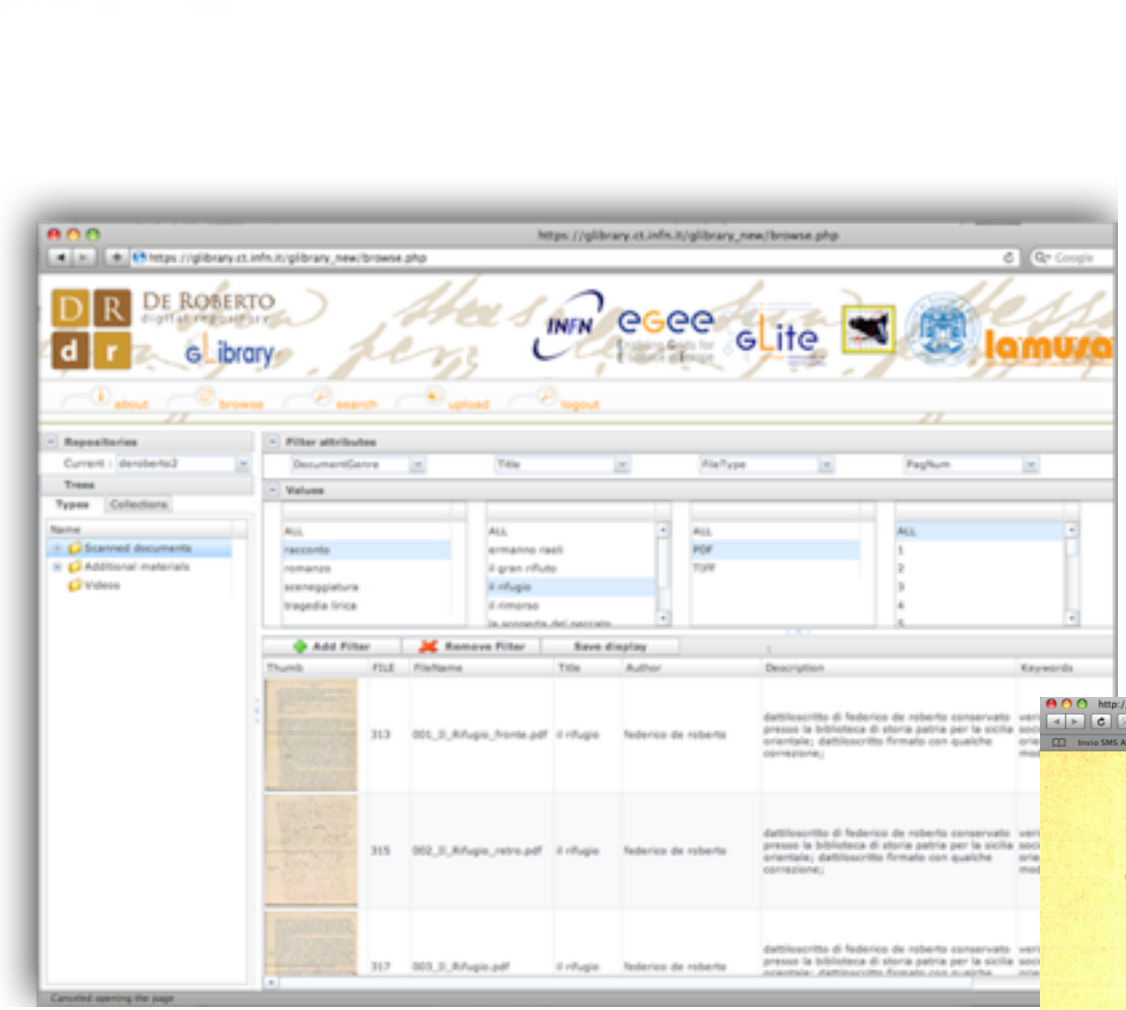

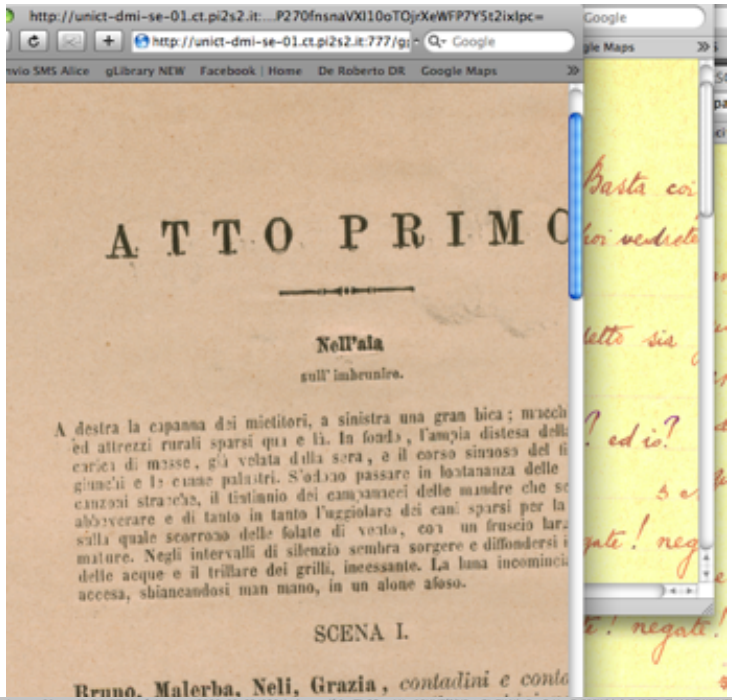

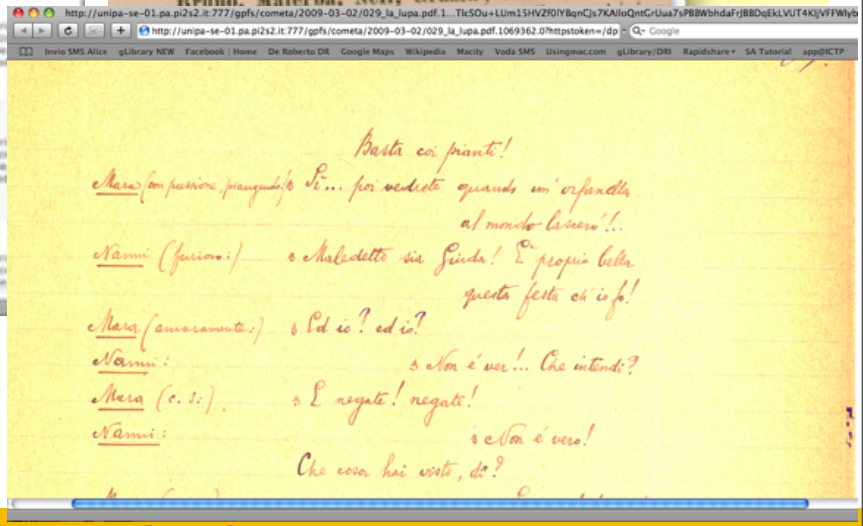

Repository dell'ESA di EO-Products

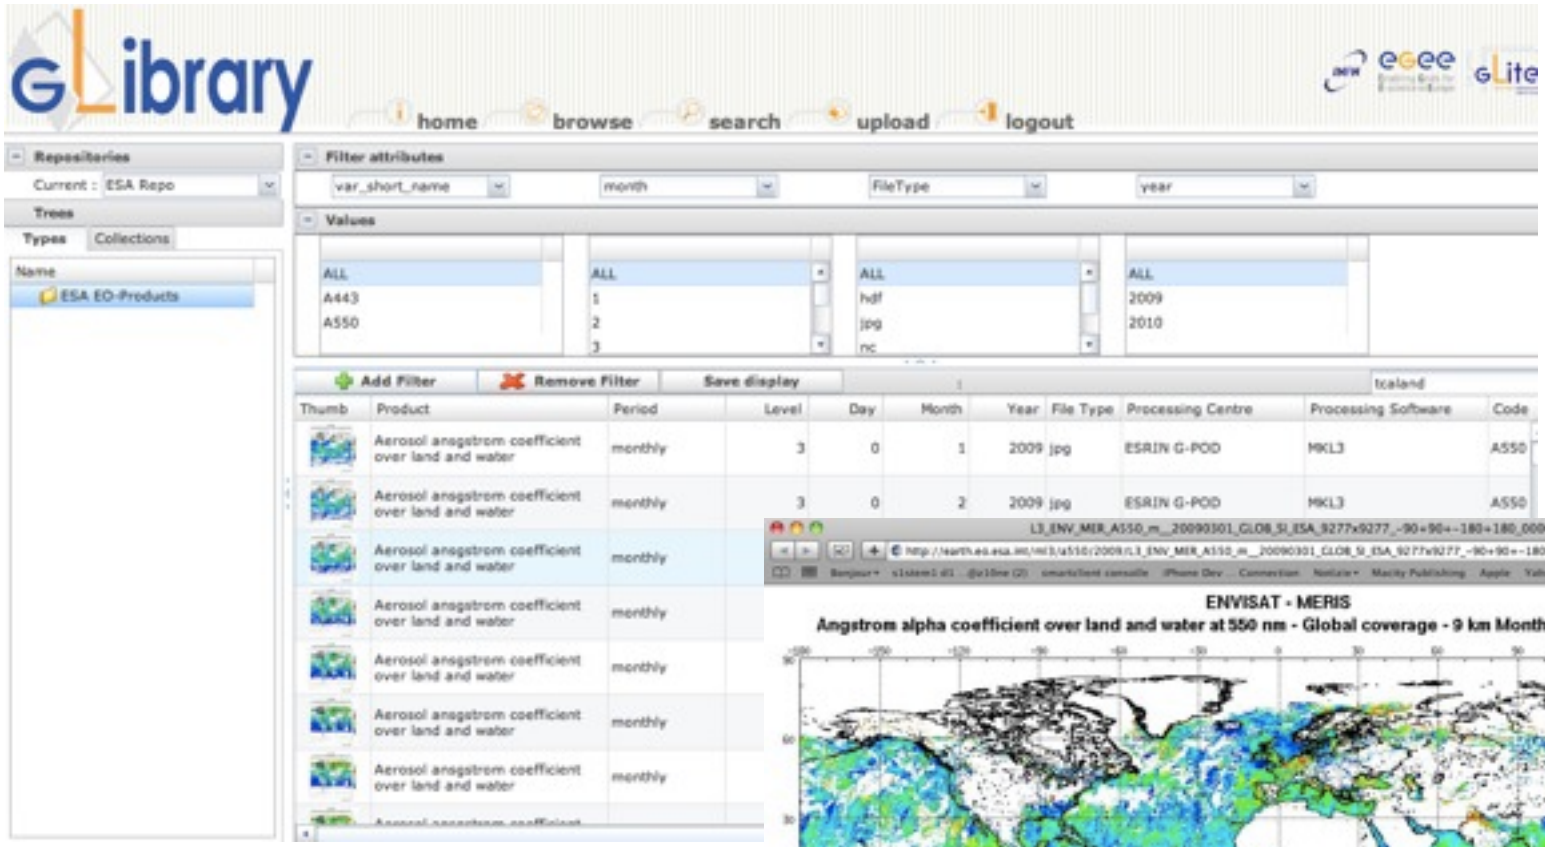

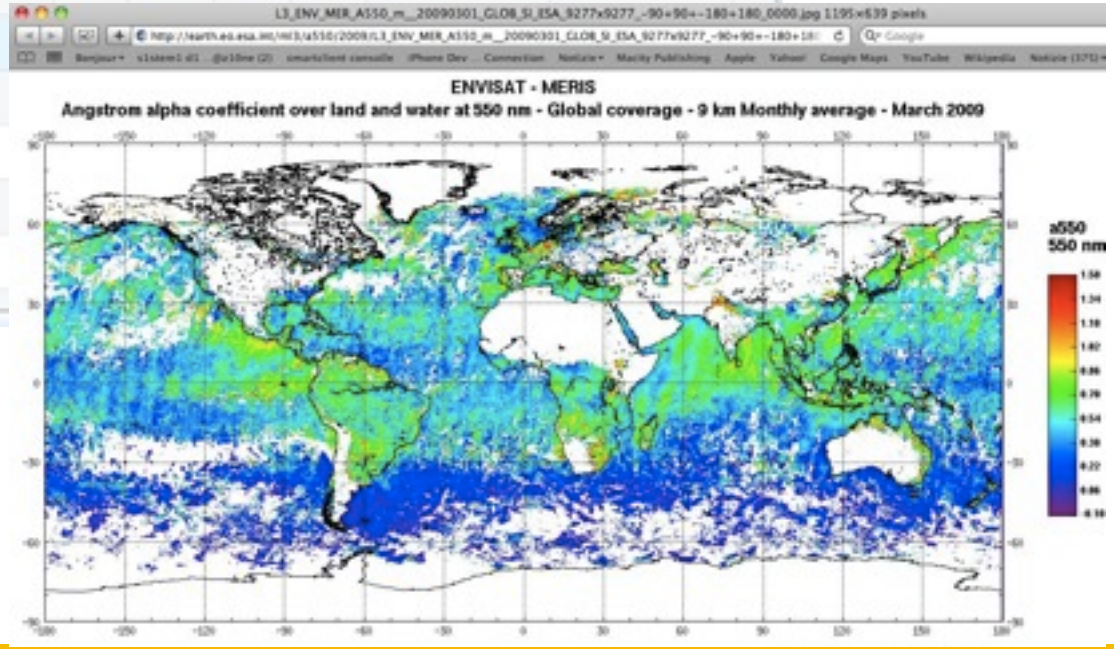

gLibrary from the iPhone

- **Demo presented at the last EGEE UF5 in Uppsala:**

- **Some screenshots here:**
  - https://glibrary.ct.infn.it/iphone/screenshots/glibrary_mobile_screenshots.pdf

- **and a YouTube video here:**
  - http://www.youtube.com/watch?v=jl4c8ZjdER4

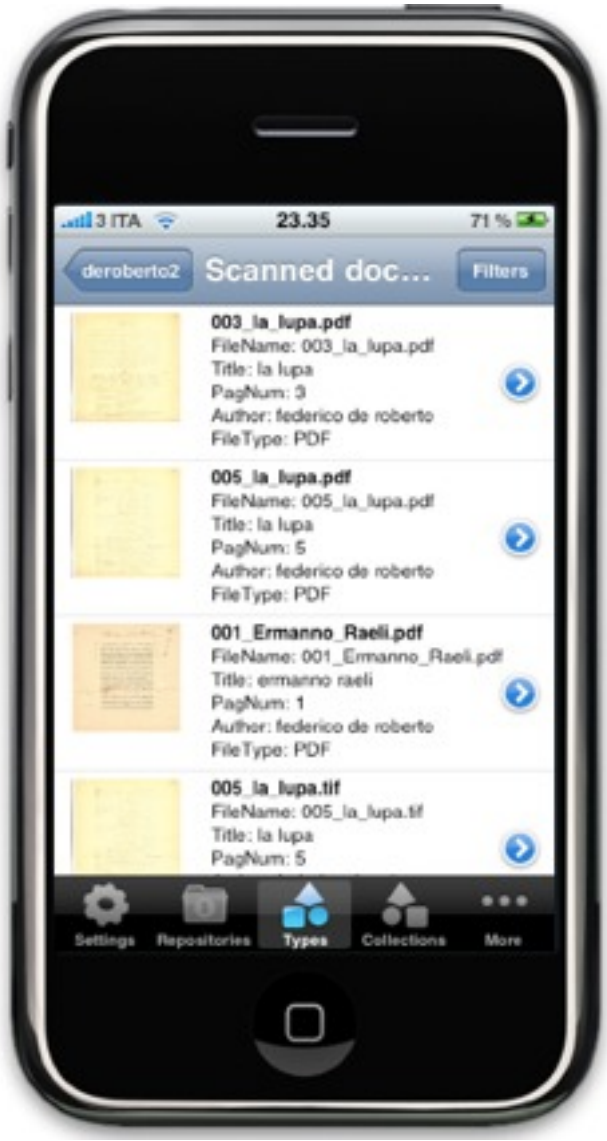

gLibrary sull' iPhone

- **Vantaggi:**
  - fornisce un'interfaccia utente estremamente **intuitiva** e basata sul **tocco** per accedere ai metadati e allo storage distribuito
    - consente l'accesso a enormi capacità di storage a utenti non esperti di grid
  - (to do) selezione automatica della replica più vicina, ottenendo le locazione corrente dal GPS
  - permette la consultazione off-line degli assets precedentemente scaricati

References

- **AMGA Web Site**
  http://cern.ch/amga
- **AMGA Manual**
  http://amga.web.cern.ch/amga/downloads/amga-manual_1_3_0.pdf
- **AMGA API Javadoc**
  http://amga.web.cern.ch/amga/javadoc/index.html
- **AMGA Web Frontend**
  http://gilda-forge.ct.infn.it/projects/amgawi/
- **AMGA Basic Tutorial**
  https://grid.ct.infn.it/twiki/bin/view/GILDA/AMGAHandsOn
- **More information on existing DB access @:**
  - http://amga.web.cern.ch/amga/importing.html
  - https://grid.ct.infn.it/twiki/bin/view/GILDA/AMGADBaccess

gLibrary references

- **gLibray project homepage:**
  - https://glibrary.ct.infn.it/
- **gLibrary paper:**
  - https://glibrary.ct.infn.it/glibrary/downloads/gLibrary_paper_v2.pdf

Questions…