

AMGA Metadata Server

- Metadata Services in gLite

(+ ARDA DB Deployment Plans with Experiments)

Birger Koblitz

NA4 3D Workshop, Oct 17th 2005, CERN

- Metadata services on the Grid comes in 2 flavours:
 - File metadata

Files	
LFN	Production

- Simple, generalized rel. DB services:

Images		
GUID	Date	Patient


Patient	
ID	Doctor

Doctor	
Name	Hospital


Example from
EGEE-BioMed community

AMGA is the Metadata Catalogue for gLite:


- **AMGA started out as ARDA's tool to investigate metadata access on the GRID**
- **AMGA will be in gLite release 1.5**
- **AMGA works in 2 modes:**
 - Side-by-Side a File Catalogue (LFC): File Metadata
 - Standalone: General relational data on Grid
- **AMGA has 2 front ends:**
 - SOAP with PTF standardised interface
 - Text-based TCP streaming protocol (proprietary, documented)
- **AMGA has ideas from many people:**
UK GridPP Metadata Group, GAG (HEP),
gLite DM-team, PTF, LHCb
- <http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/>




- **AMGA implements a common interface designed in close collaboration of gLite and ARDA teams**
(P. Kunszt, R. Rocha, N. Santos, B.K.)
- **Again: many ideas from UK GridPP Metadata group, LHCb (Bookkeeping, GANGA), GAG, PTF...**
- **Endorsed by PTF**
- **Design Ideas:**
 - Versatility: Usable for HEP as well as Biomed (security)
 - Modular: Interface for Entry manipulation, schemes, security
 - Possible Add-on to File Catalogue
 - Allows stateless & statefull implementations
 - Few requirements on back end, can be €
- **Description of WSDL at**
<https://edms.cern.ch/document/573725>



- Traditional DB access doesn't work on Grid:

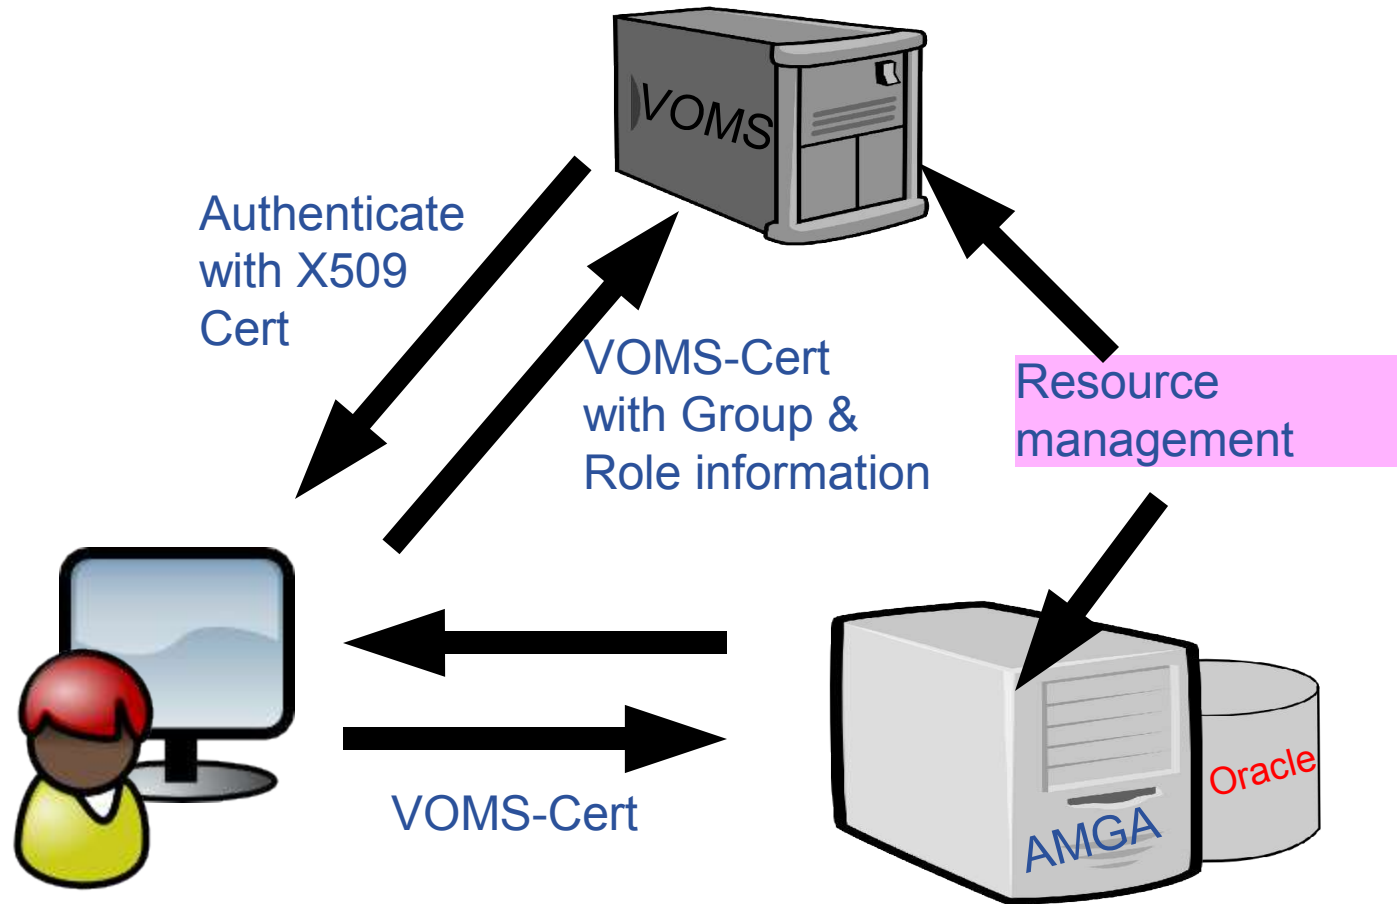


- +Performance
- +Simple Implementation
- Security, Monitoring
- Authentication, resource management??



- +Lightweight Client
- +Security: GSI, x509
- Performance
- Implementation: State

- Access control to resources on the Grid is done via a Virtual Organization Management System (VOMS):



- **Security very important for BioMed, less for HEP**

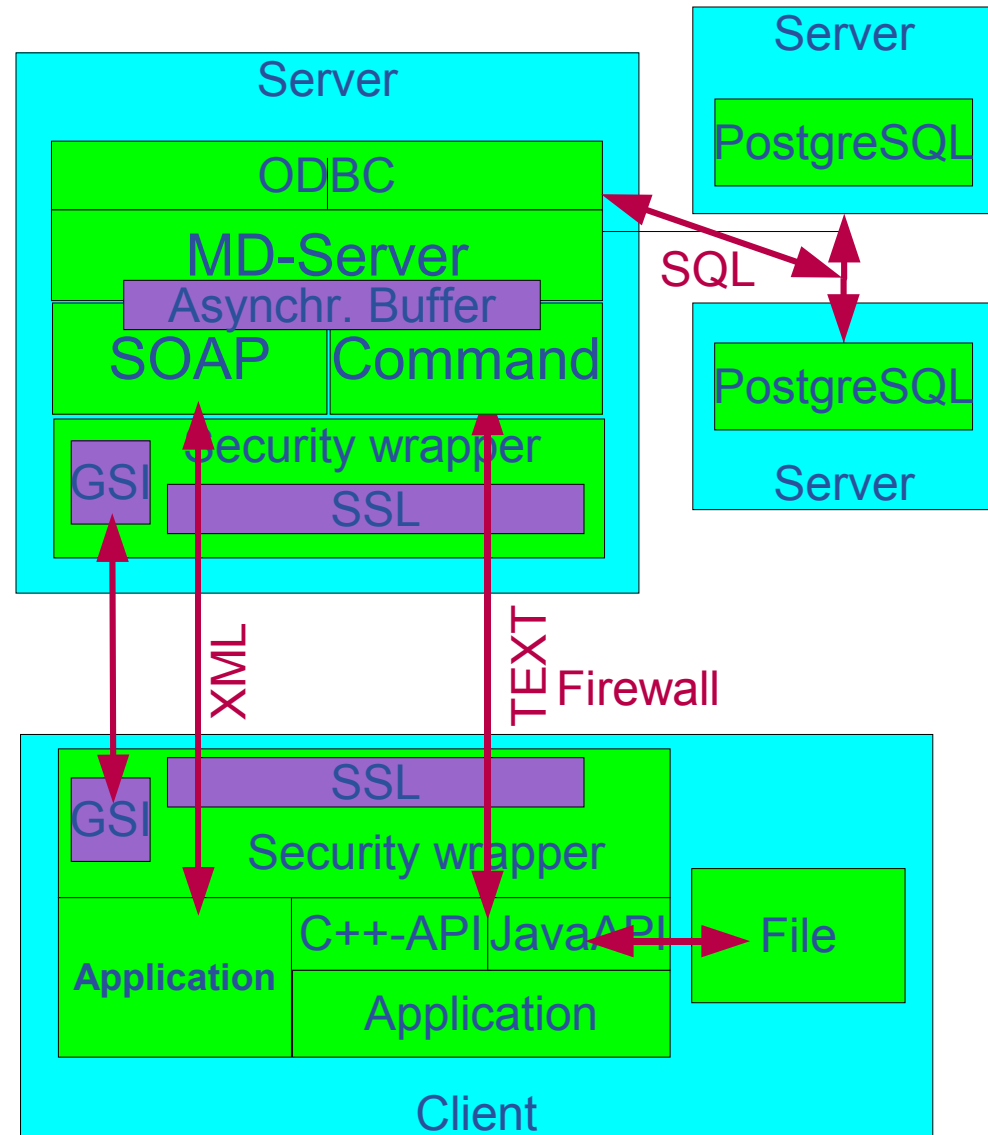
Security ↔ Speed

- **Standalone catalogue has:**
 - ACLs for dirs and Unix permissions dirs/entries
 - Built-in group-management as in AFS
 - **AMGA + LFC back end:**
 - Posix ACLs + Unix permissions for dirs/entries (ACLs currently not checked: slow!)
 - Users/groups via VOMS
 - **Currently no security on attribute basis**
 - AMGA allows to create **views**: Safer, faster, similar to RDBMS
- Security tested by GILDA team for standalone catalogue, liked built-in group management & ACLs, but we **need feedback from BioMed!**

AMGA in preproduction within several projects:


- **LHCb and ATLAS: GANGA**
- **LHCb Logging and Bookkeeping**
- **EGEE BioMed applications**
 - Highly secure access to medical images metadata
- **Generic applications:**
 - Metadata for EGEE-GILDA media library
 - UNOSAT project: Used side-by side with LFC catalogue for file-metadata of satellite images

- **AMGA Implementation:**
 - SOAP and Text frontends
 - Supports single calls, sessions & connections
 - SSL security with grid certs
 - PostgreSQL, Oracle, MySQL, SQLite backends
 - Works alongside LFC
 - C++, Java, Python clients
- **See & download at**
<http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/>



PhD (N. Santos) focuses on replication of AMGA at application level:

- **Goal**
 - Fault-tolerance, scalability, reliability with minimum administration overhead
- **Ideas**
 - Replication.
 - Partial, only collections that are needed locally.
 - Master-slave asynchronous
 - P2P techniques for managing dynamic network of nodes
 - Discovery and location of metadata collections using P2P techniques
- **Current status:**
 - Initial work on AMGA
 - Prototype using PostgreSQL replication working



ARDA Database Deployment Plans and Replication Test Plans

Analysis Prototypes developed by HEP-Experiments & ARDA currently foresee:


- **LHCb:**
 - Logging & Bookkeeping DB replicated $\geq 2x$
AMGA + DB: Oracle
Could use Oracle streams or AMGA native replication
 - GANGA analysis front-end plans to deploy AMGA+Postgres back-end once per Tier1
Replication scheme unclear, data not persistent
- **ATLAS:**
 - GANGA analysis front-end using AMGA, DB back-end unknown (currently PostgreSQL)
- **CMS:**
 - Central dashboard DB at CERN for monitoring data aggregation
Currently PostgreSQL, PHP-access

- Generic AMGA for BioMed and other communities will be deployed by EGEE deployment teams

Rough overview with numbers:

	Replication Scheme	Update Rate	Size
GANGA (LHCb)	1DB per Tier 1	users x jobs	
GANGA(ATLAS)	?	users x jobs	
Bookkeeping(LHCb)	1+1 at Cern	1 x job	15 GB/year
Dashboard(CMS)	1+1 at Cern	10 x jobs: 1/s	30 GB/year

- ARDA would like to test Oracle streams replication in simple setup at CERN for AMGA:
 - Test replication functionality, in particular to replicate schema changes
 - Understand performance requirements



➔ **Verify Oracle Streams as replication solution for AMGA**

- **With AMGA gLite will get generic metadata service**
- **Implements most of PTF Metadata Interface**
- **AMGA has seen heavy performance/stability testing**
- **AMGA currently evaluated / in preproduction for LHCb(GANGA, bookkeeping), BioMed, UNOSAT, ESR**

- **ARDA would like to test AMGA replication via Oracle streams**
- **Replication of AMGA at application level planned**

- **Entry**
 - Has **key** (unique string) and **attributes**
- **Attribute**
 - Has **name** (string),
type (depends on backend, support for basic types)
 - Belongs to **schema**
 - An entry in a schema **has a value** for each attribute
- **Schema (in AMGA: directory)**
 - Has **name** and **list of attributes**
 - In AMGA: Every entry belongs to one schema, schemas are hierarchical: /collaboration1/jobs
- **Query**
 - SELECT ... WHERE ... clause in **SQL-like query language**

Example command line session with LFC back end:

```
mdclient -p8822 lxb0709
Connected to lxb0709:8822
ARDA Metadata Server 0.9.4
Query> dir /
>> grid
>> collection
Query> dir /grid/arda
>> lfn-0.dat
    [... rest of LFC entries]
Query> schema_create /grid/arda i int t text
Query> listattr /grid/arda
>> i
>> int
>> t
>> text
```

```
Query> addentries /grid/arda/lfn-0.dat /grid/arda/lfn-1.dat
Query> listentries /grid//arda
>> lfn-0.dat
>> lfn-1.dat
Query> addentry /grid/arda/lfn-4.dat i 2 t 'A test'
Query> listentries /grid/arda
>> lfn-0.dat
>> lfn-1.dat
>> lfn-2.dat
Query> addattr /grid/arda f float
Query> find /grid/arda/* 'i=2'
>> lfn-2.dat
Query> addentries /grid/dteam/arda/bla.dat
Error 1: File or directory not found: /grid/arda/bla.dat
Query>
```