

Replicating Metadata with AMGA

B. Koblitz, CERN

with N. Santos and V.Pose

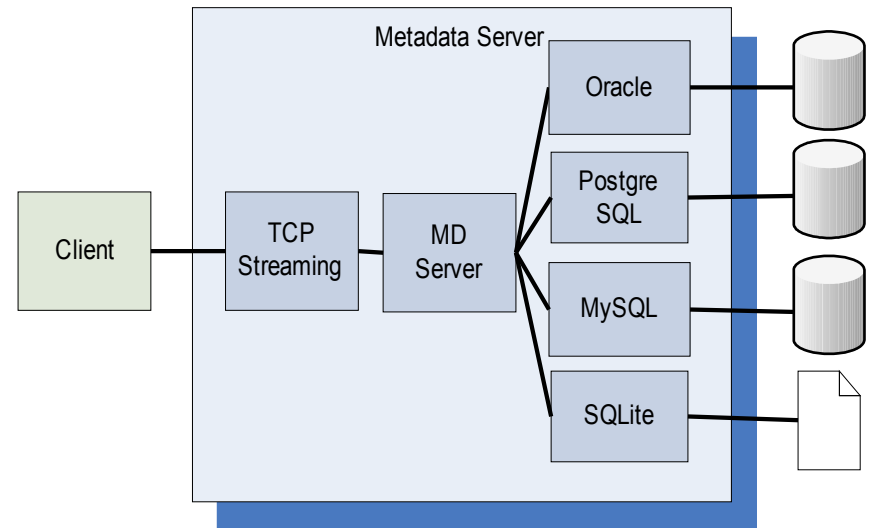
GridPP Metadata Workshop Oxford

July 5th, 2006

- **What is AMGA? What is Metadata on the Grid?**
- **Interface definition**
- **Implementation of AMGA**
- **Comparison with other Catalogue services**
 - Benchmarks of LFC, FiReMan and AMGA
- **Replication of Metadata with AMGA**

- **AMGA Implementation:**

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




- **AMGA in (pre-)production within several projects:**
- **LHCb and ATLAS: GANGA**
 - Highly dynamic relational data for tracking jobs
 - AMGA used for communication between user clients
- **LHCb Logging and Bookkeeping**
 - 15 Million log entries of jobs
- **EGEE BioMed Medical DataManger**
 - Highly secure access to medical images metadata
 - Many different security roles, uses VOMS
- **Generic applications:**
 - Metadata for EGEE-GILDA DigitalLibrary
 - UNOSAT project: Satellite images metadata, GIS

AMGA-Replication implemented as PhD project (N. Santos):


- **Scalability:** More connections, smaller latencies
- **Reliability:** Avoid single point of failures
- **Administration:** Less DBA time central management possible
- **AMGA implements replication & federation itself:**
 - Asynchronous log-shipping of Metadata commands: **cross-vendor**
 - **Users can changes replication and federation setup**
- **Initial implementation currently tested**
- **Federation capabilities essential for EGEE Medical DataManager**
 - Keep sensitive patient information distributed

Database Replication for LHC:



EGEE Grids (e.g. BioMed):



- Tiered Architecture
- Replication with Oracle Streams to Tier 1
- Require Oracle on T1





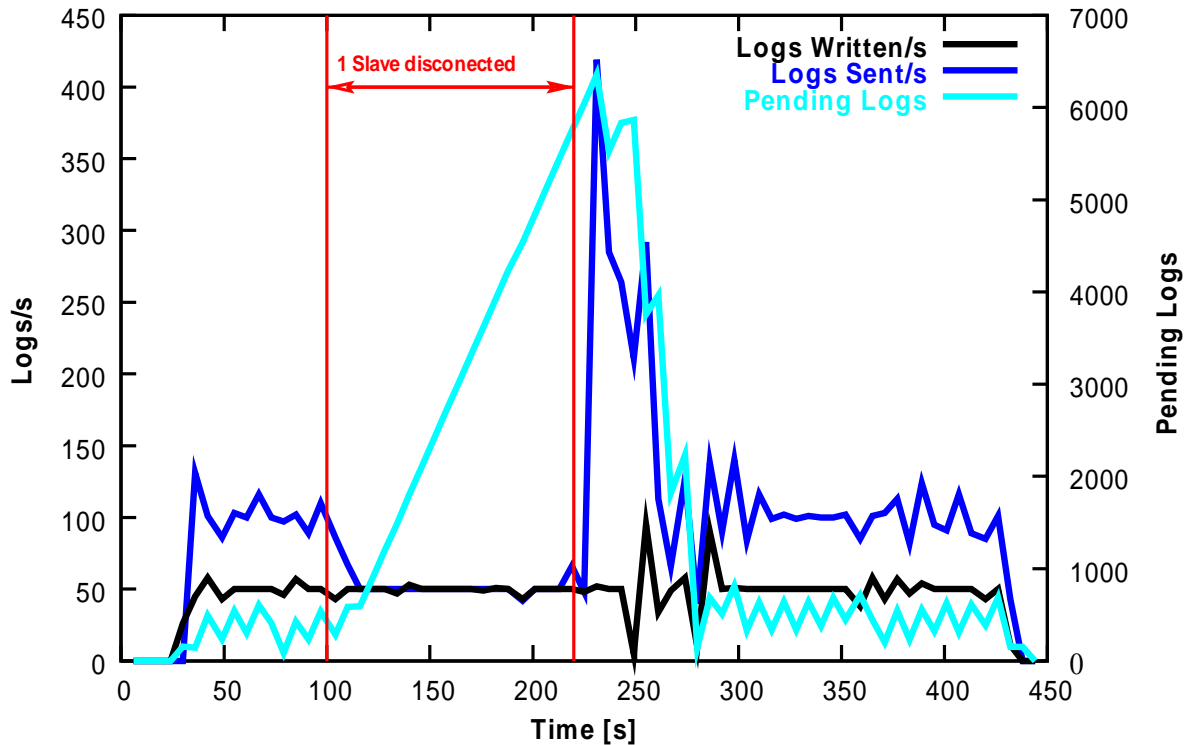
- Loose cluster of sites
- Diverse infrastructure (MySQL, Oracle, ...)




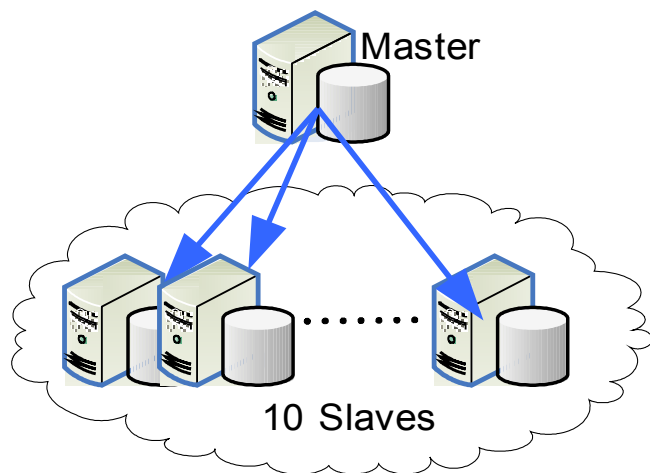
- **Database replication:**
 - Vendor dependent
 - Replicates nearly all application data
- **3D Project (IT-PSS):**
 - CERN -> T1 replication with Oracle Streams
- **Collaboration with 3D team on Monitoring & Benchmarks**
- **Metadata replication:**
 - DB independent, cross DB
 - Complex implementation
 - Customized to Metadata

- AMGA replication makes use of **hierarchical concept**:



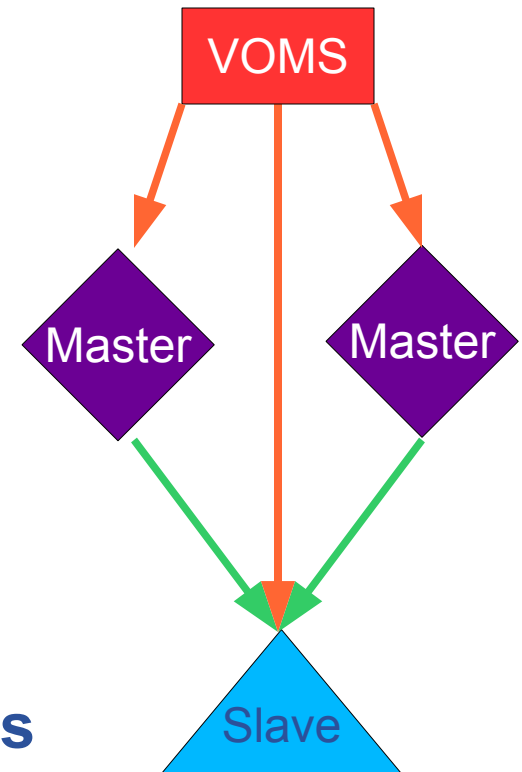
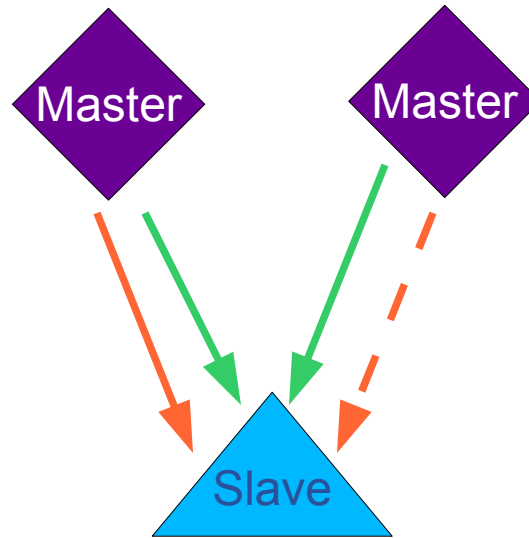



- **AMGA keeps logs for disconnected slaves**
- **Reconnected slaves are brought up-to-date automatically**
- **Fast recovery**
- **Need long term tests -> late Summer!**



- **Scalability test**

- Setup with 1-10 slaves, master is desktop PC (3GHz P4)
- Inserts done with a single client at maximum client speed (90 entries/s)
- Replication overhead: 20% more CPU
- System scales very well: 10% CPU-Usage for 10 slaves



- **AMGA shall support 3 Replication modes**
 - No ACLs, no user information
 - ACLs, user information from VOMS
 - ACLs and user information

- **AMGA implements very flexible replication and federation features**
- **Provides Cross-Vendor replication**
 - Oracle -> MySQL
- **For HEP replication replication is very interesting because of tiered structure**
- **Initial tests show good stability and performance**
- **Next topic is replicated User Management**