Information and Monitoring Status and Plans

JRA1 All-Hands, Helsinki, 18-20 June, 2007

Steve Fisher/RAL on behalf of JRA1-UK

www.eu-egee.org
www.glite.org

Information Society and Media

Status

- **R-GMA**
  - Some recent problems with the registry being overloaded
    - See later
  - Coding and testing of parts of new design almost complete
    - See later
- **bdII**
  - Indexing added
- **SD**
  - Stable
  - SAGA code being developed
    - See later

Work-plan for GLUE etc.

- **Version 1.3 of the Glue Schema (CERN, CCLRC)**
  - LDAP schema in certification (backwards compatible)
  - R-GMA schema will be updated to match
  - *Note that GLUE 2 will be defined by new OGF working group*
- **Write GIP Info Provider for Services (CCLRC)**
  - When the schema is in production
  - Being done by Stephen Burke

Current Registry Problem

- **The load on the registry (about 30 requests a second) was too great for the current implementation**
- **Short term solutions:**
  - Reduce load by having one GIN per MON box
    - Patch 1093 is marked as high priority
  - Reduce load by requesting that SAM job wrapper no longer publishes to R-GMA by default
    - Not under our control
  - Registry improvements (deployed last Friday):
    - Better locking on registry
    - Connection pooling to MySQL
    - More indexing on MySQL
    - Optimised buffer sizes in MySQL
- **Registry can now cope - better with Patch1093**

Registry – longer term solution

- **More registry instances**
  - Multiple VDBs
  - Registry replication
  - Ready towards end of the year

First release of new R-GMA design - I

- **Redesigned Schema and Consumer but rest unchanged**
  - Schema and Consumer (and later other components) in one Servlet
    - Makes inter-service calls on same node very fast
    - Can share some objects more easily between services on same node (e.g. TaskQueue)
    - Makes use of Listener (JDK 5) to detect memory shortage in good time
      - *We can then send an RGMABusyException in response to requests that would increase memory use*
- **Schema**
  - With replication
    - One master per VDB
    - Slaves pull updates from master ("all changes since … " to avoid queues on master
  - Multiple VDB support ready

First release of new R-GMA design - II

- **Consumer**
  - Able to stream from old producers (one connection per producer) and talk to old registry
  - Ready to stream from new producers (single connection to Mon box) and to talk to new Registry
  - For continuous queries polls registry looking for relevant producers
- **Status**
  - It builds as a single .war file with the new “big” servlet and the old servlets
  - Testing in progress
  - Modifications to the inspector being completed

Then in sequence (overlap is OK)

- **Primary Producer**
  - Able to stream to old and new Consumers
  - Only one socket for streaming from one Mon box to another when streaming to new Consumer
  - Database independence
  - Managed tuple stores - essential to support authz
  - Quite a lot already written
- **Secondary Producer**
  - One component
    - Data moved directly into tuple store without multiple translations
- **On Demand Producer**
- **Registry**
  - The registry no longer sends out notifications
    - “old” Consumers will then not work so there will be some delay after initial release on previous slide
    - Should increase reliability
    - Registry replication will be much simpler
  - Multiple VDB support but no cross VDB queries yet
  - Quite a lot written already
- **Browser**

Subsequent releases

- **Will provide in this sequence:**
  1. Queries over multiple VDBs
     - Almost standard SQL
       - *Extension for Unions*
  2. Authz by VDB
     - This will make use of certificate attributes (VOMS groups/roles)
     - Database engine is used to implement parameterised views
     - Code developed for older version of code by TCD for consumer
     - May be done first or in parallel with above
  3. Registry replication
     - Much easier now that registry is passive
  4. Oracle support
     - DB independence part of new design

Testing

- **To ensure that the new code is well tested our testing framework is being overhauled**
  - Make it easier to run even for those tests involving many machines
  - It includes simulations of various R-GMA applications that can be run in parallel for a long period

Service Discovery

- **SD API Specification presented at OGF20 (Manchester)**
  - Some changes suggested
  - Document will be resubmitted to SAGA group soon
- **Implementation**
  - Currently
    - bdII C++ plug-in is working
    - R-GMA C++ plug-in will be done soon
    - wrapper to look like existing gLite SD completed
  - Next
    - Modify to match final SAGA specification
    - release
    - Implement SAGA plug-in mechanism
      - *Want to allow multiple back-ends simultaneously*
    - C and python implementation as wrappers
    - Java implementation

Service Discovery (SD) Plans

- **"Configuration-free" SD (CCLRC)**
  - Useful as a bootstrap mechanism
    - it can locate the information server on the local subnet
  - Will use an existing protocol
- **Make use of the SD APIs in all components (All!)**

GLUE 1.3 Service

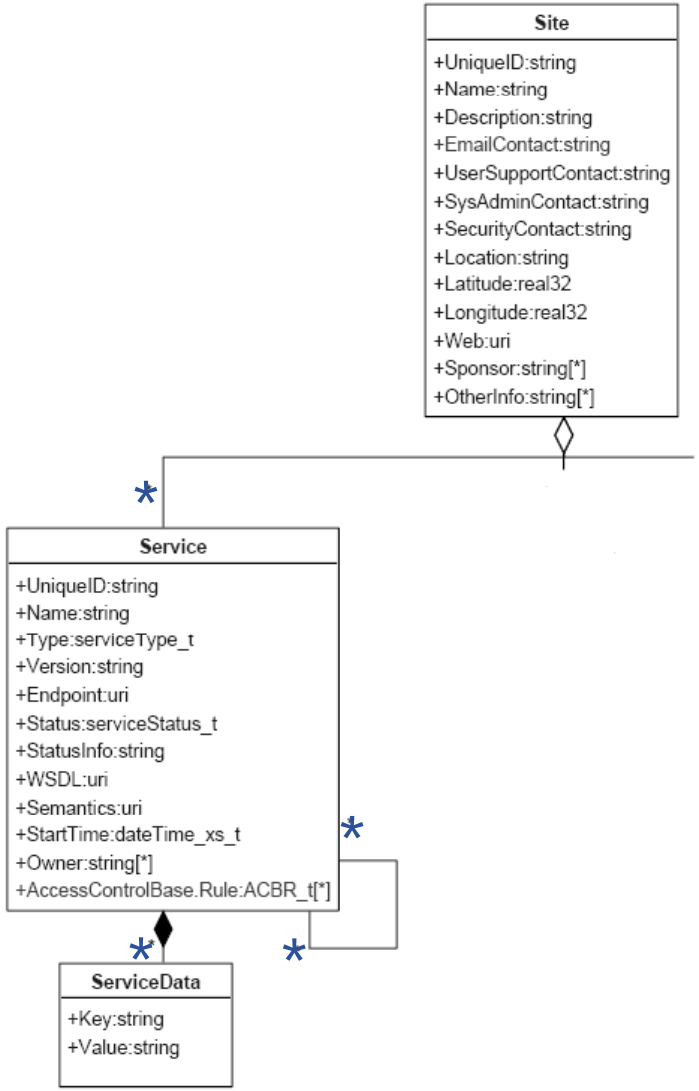

- **Site may have many services**
- **Services have n:n self-relationship**
- **Service may have service data**
  - (key, value)

GLUE 2.0 Service

- **Being developed by new OGF WG with co-chairs:**
  - Sergio Andreozzi
  - Balazs Konya
  - Laurence Field
- **I am responsible for the part which is neither CE nor SE i.e. Service, Resource, Site …**
- **Some of co-chairs also part of OGSA-WG Info Services "Design Team"**
  - Communication is happening
- **GLUE 2 will be different**
  - Not much else to say yet

INFOD-WG

- **Very R-GMA like but much more general**
- **Being noticed by OGSA-WG**
- **Implementation**
  - M.Sc. Student at Edinburgh
    - Simple incomplete prototype using open source components
    - Compare it to other systems
    - Feedback on gaps/errors in the specification
  - University of Tennessee led collaboration
    - To produce and deploy an open source implementation
    - Weekly phone meetings
    - Team of 8
    - Progressing well

Mailing lists and web

- **R-GMA-ANNOUNCE**
  - Low volume for R-GMA announcements to users
  - Moderated
  - Replies go to R-GMA-SUPPORT list
- **R-GMA-SUPPORT – r-gma-support@physics.gla.ac.uk**
- **R-GMA-DISCUSS – r-gma-discuss@physics.gla.ac.uk**

- **To subscribe to the announce or discuss list:**
  - http://www.physics.gla.ac.uk/mailman/listinfo/<list name>

- **SD lists will be created**

- **Web: http://hepunx.rl.ac.uk/egee/jra1-uk/**