

# Information and Monitoring Status and Plans

*JRA1 All-Hands, Catania, 7-9 Mar, 2007*

*Steve Fisher/RAL on behalf of JRA1-UK*

[www.eu-egee.org](http://www.eu-egee.org)  
[www.glite.org](http://www.glite.org)

- **People**
  - All posts filled
    - Parminder Bhatti joined us since Cosener's meeting
- **R-GMA**
  - Stable - but some problems with heavily used Mon boxes
    - We think we understand it and have a fix
  - Coding and testing of parts of new design almost complete
    - See later
- **BDII**
  - stable
- **SD**
  - Stable
  - SAGA style code being developed
    - See later

- **Version 1.3 of the Glue Schema (CERN, CCLRC)**
  - November
  - 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 (CERN)**
  - September
  - Will be done once the schema is in production
- **Investigate Debian system to tar up logs and make them available (CCLRC, CERN)**
  - *October – and still slipping*

- **Scripts to monitor an R-GMA installation (CCLRC)**
  - November
  - Done

- **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
- **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

- **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
- **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
- **Browser**

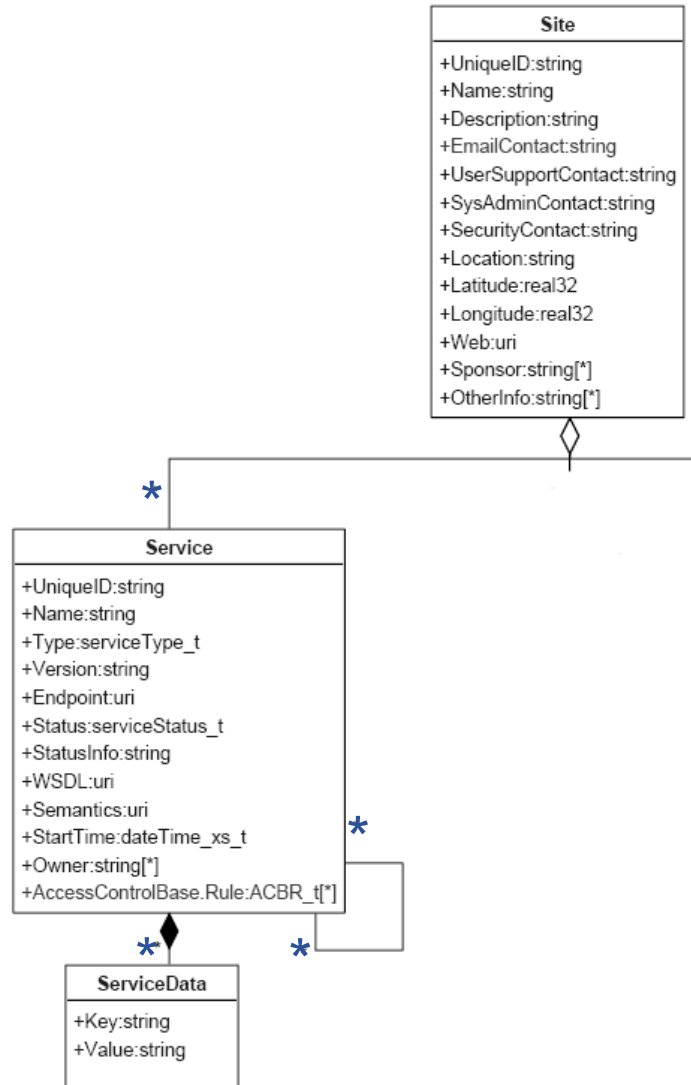
- **Will provide in this sequence:**
  1. Queries over multiple VDBs
    - Almost standard SQL
      - *Extension for Unions*
  2. Authz by VDB
    - This will make use of certificate attributes (VOMS groups/roles)
    - Database engine is used to implement parameterised views
  3. Registry replication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



- **Modify mediator to make use of secondary producers with a predicate**
  - This allows data to be partitioned and thereby reduces the total amount of data processed by one secondary producer
- **Allow a tree of secondary producers**
  - This is a generalisation of the above
  - To be effective it requires that the top level node is rarely used for queries
- **More complex predicates**
  - Currently just simple restrictions on one attribute

- **Parallel invocation of plug-ins (CCLRC)**
  - Mar 2007
  - The fastest plug-in wins
  - This should be configurable (or file wins every time!)
- **Resolve caching issues (CCLRC)**
  - Mar 2007
- **"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!)**

- **Between**
  - EGEE, OSG, OMII (Europe and UK), Nordugrid, ARC, Globus and FSU
- **Common Service Discovery interface is needed.**
  - There is ongoing work in SAGA activity within OGF.
  - A plugin specification is also being defined to enable the APIs to be used within multiple systems.
  - Similar plugins developed as part of the OGF gin-info activity.



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

- **Finding Services**
  - Based on various search criteria
  - Includes key/value pairs (open-ended)
  - Can use multiple plugins (and combine the results)
- **Returns a service “object”**
  - Has getter methods
    - Hide implementation
    - Allow changes
    - Optimal efficiency

## ListServices

IN ServiceFilterString

IN VOFilterString

IN DataFilterString

OUT List of service "objects"

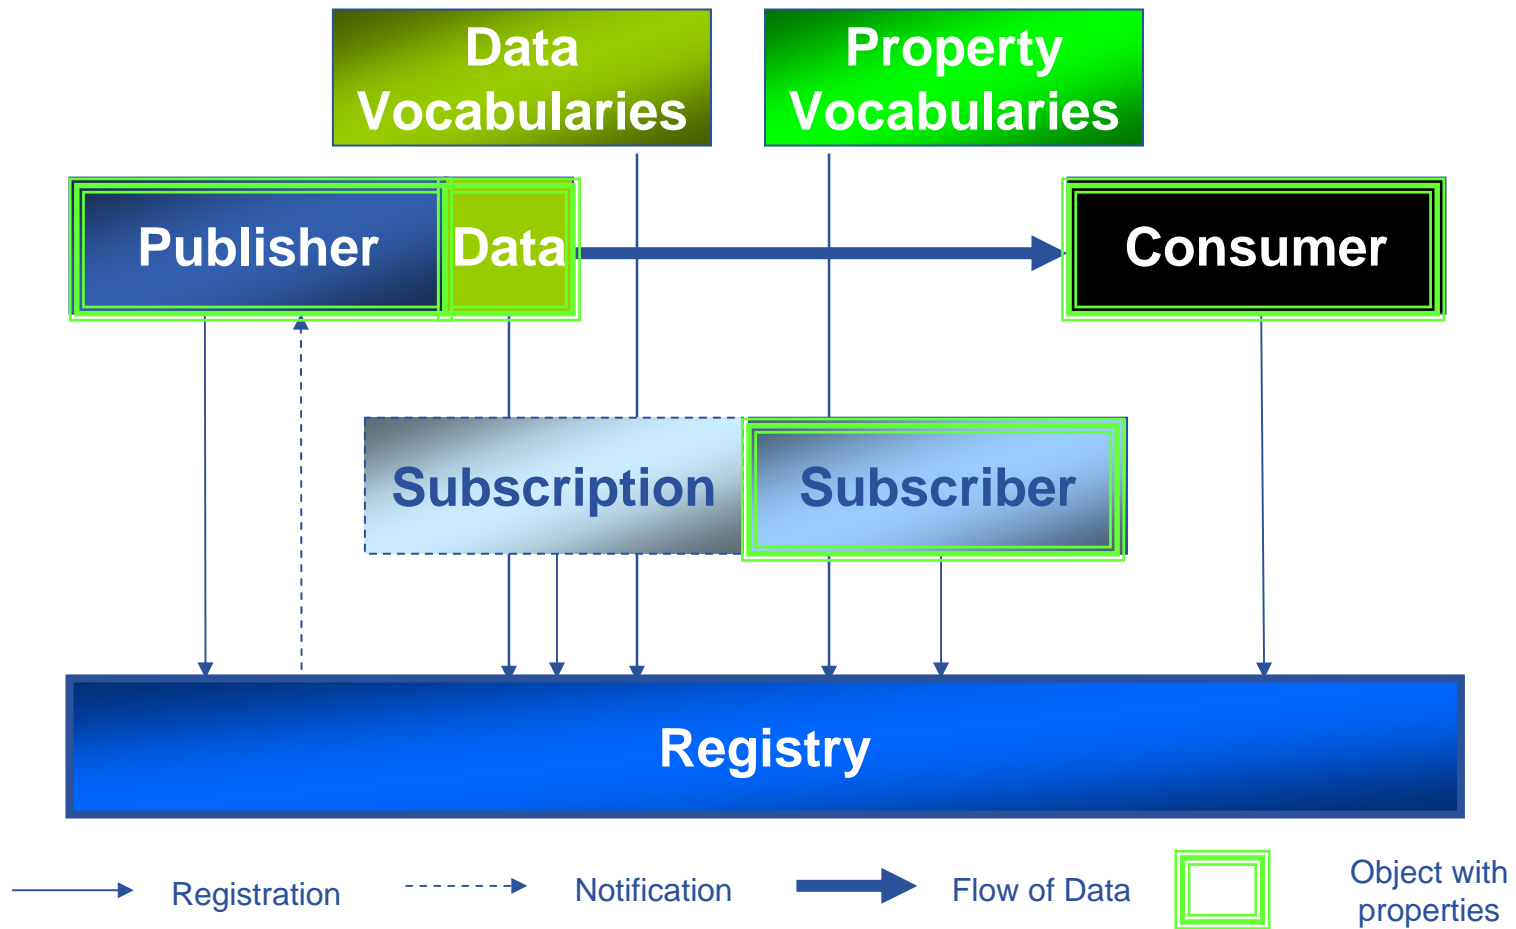
- **Filter strings uses SQL syntax as if it were part of a where clause selecting from a single table.**
- **3 filter strings**
  - simplifies the implementation,
  - clarifies the description of the functionality
  - avoids clash with key name being glue attributes.

- **Column names in the service filter are:**
  - Type - type of service
  - Name - name of service
  - Site - name of site
  - Endpoint - will normally be used with the LIKE operator
  - Service - for associated services
- **Column names in the VOFilterString are**
  - VO - will often be used with the IN operator
- **Column names in the The DataFilterString**
  - are taken from the service data key/value pairs.

- **listServices ("Site IN ('INFN-CNAF', 'RAL-LCG2') ", NULL, NULL)**
  - all services running at any of the two sites
- **listServices("Type = 'ResourceBroker' AND Site LIKE '%INFN%' ", NULL, NULL)**
  - all services matching a type and site name by pattern
- **listServices (NULL, "VO IN ('cms', 'atlas') ", NULL)**
  - all services for matching VOs
- **listServices ("Type = 'ResourceBroker' ", NULL, "RunningJobs >=1 AND RunningJobs <= 5 ")**
  - all service matching service type and key/value interval
- **listServices ("Endpoint LIKE '%PrimaryProducer%' ", NULL, NULL)**
  - all services matching end point pattern

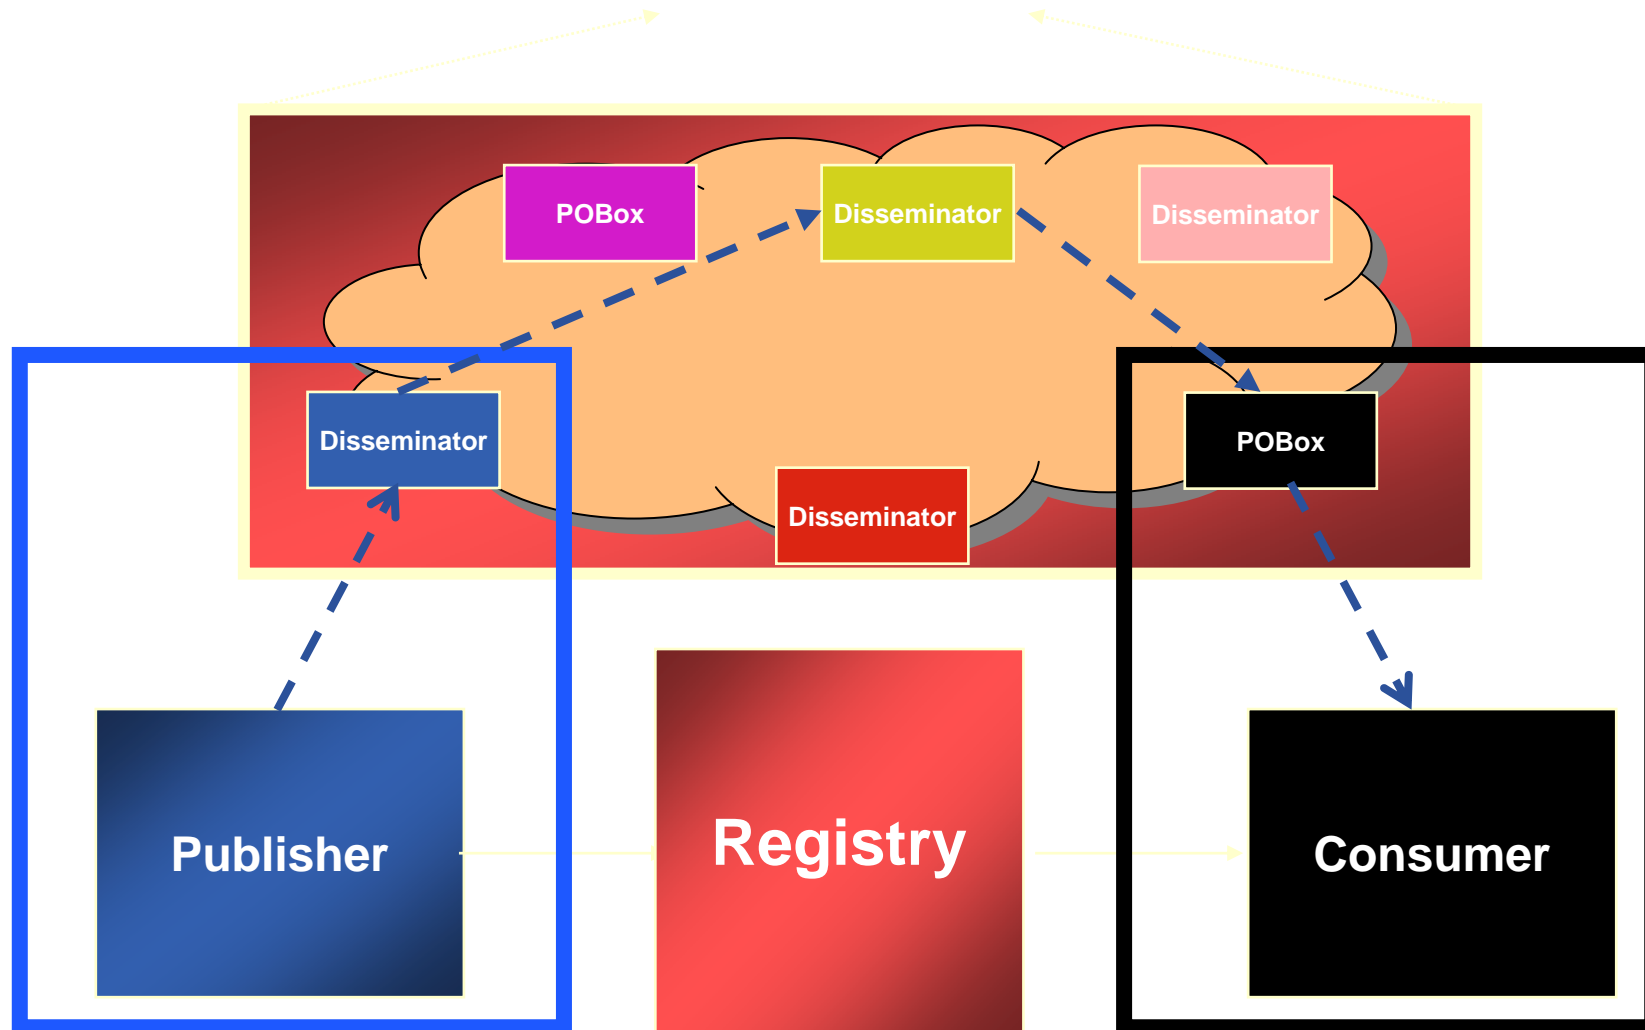


- **SAGA people have asked me to take this forward and produce Spec for user API for GGF20**
- **Work has started**
  - Have so far done nothing about conformance to SAGA style
- **Also need the plugin API so that plugin providers can start work**
  - Implementations of plugins also within gin-info
- **Currently writing a prototype 3 string selection (to replace current large set of calls)**
  - Have sorted out SQL parser to generate R-GMA calls (trivial) and good progress with LDAP one
  - Using ANTLR rather than JavaCC because of C++ support
- **Will provide a compatibility layer on top**



- **60 day comment period ended**
- **9 comments were received**
  - Responses produced as appropriate
    - A lot of effort went into these responses
- **Have almost finished making changes**
  - Avoid “new features”
    - Want to achieve acceptance ASAP
- **Will learn more from the implementations**

- **M.Sc. Student at Edinburgh**
  - Simple incomplete prototype using open source components
  - Compare it to other systems
  - Feedback on gaps/errors in the specification
- **Oracle**
  - Intends to apply INFOD technology in customer cases. This could lead to a reference implementation. Oracle does not comment on future product plans
- **Collaboration**
  - University of Tennessee, Oak Ridge, IBM and Oracle
  - To produce and deploy an open source implementation
  - This seems to be going well



- **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](mailto:r-gma-support@physics.gla.ac.uk)
- **R-GMA-DISCUSS** – [r-gma-discuss@physics.gla.ac.uk](mailto: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>>
  
- **Web:** <http://hepunx.rl.ac.uk/egee/jra1-uk/>