Information and Monitoring Status and Plans

JRA1 All-Hands, CERN, 24-26 Oct 2007

Steve Fisher/RAL on behalf of JRA1-UK

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

Information Society and Media

Status

- **R-GMA**
  - Stable
  - Coding and testing of most of new design almost complete
    - See later
- **bdII**
  - Stable
- **SD**
  - Stable
  - SAGA code being developed
    - See later

Work-plan for GLUE etc.

- **Version 1.3 of the Glue Schema (CERN, STFC)**
  - LDAP schema (patch 980) in production
    - DONE
  - R-GMA schema will be updated to match
    - Not done yet
- **Write GIP Info Provider for Services (STFC)**
  - DONE (Stephen Burke)

First release of new R-GMA design

- **This includes a lot more than originally planned**
  - Better use of developer's time
  - Easier to test new Consumer after new Producer written

- **Includes:**
  - Services:
    - Primary Producer
    - Consumer
    - Schema
    - Registry
    - Inspector
  - Facilities
    - Single servlet
    - Task Manager
    - Tuple Store
    - Multiple VDBs
    - Authz
    - Replication of schema and registry
    - Improved streaming

- **Excludes (currently)**
  - Services
    - Secondary Producer
    - On Demand Producer
    - Browser
  - Facilities
    - Oracle support

Facilities - I

- **All refactored 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 send an RGMABusyException in response to requests that would increase memory use
- **Task Manager**
  - Most communication between services is put on a task queue
    - Allows regular handling of tasks that do not proceed as planned
  - Tasks have an associated key
    - Only tasks with a “good” key may run in parallel
    - Definition of goodness is dynamic
- **Improved streaming**
  - Makes use of NIO channels to reduce connections between machines

Facilities - II

- **Tuple Store**
  - Defined by interface
  - Multiple implementations:
    - HSQLDB
    - MySQL
    - Oracle to be written
  - Authz enforced at tuple store
- **Authz**
  - Fine grained for R and W
  - Based on parameterised views stored in schema
  - Code provided by Stuart Kenny – TCD
- **Multiple VDBs**
  - Provide separate name spaces (own schema and registry)
  - Essential for scaling
- **Replication of schema and registry**
  - Very different algorithms
    - Master schema – slaves ask for changes since
    - Registry pushes changes to peers

Probably not in first release

- **Secondary Producer**
  - One component
  - Data moved directly into tuple store
- **On Demand Producer**
- **Browser**
- **Oracle support**

Testing

- **Because of single servlet and a piece of code to detect local communication can now make end to end tests without deployment as simple unit tests**
- **Our testing framework is also 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 Spec re-presented in SAGA at OGF21 (Seattle)**
  - With very minor change we are encouraged to submit for public comment
  - N.B SAGA are also planning to tackle IS in general
- **Implementation**
  - Currently
    - bdII and R-GMA C++ plug-in is working
    - Wrapper to look like existing gLite SD completed
    - Code given to Ales to try
  - Next
    - C and python implementation as wrappers
    - Java implementation

Service Discovery (SD) Plans

- **"Configuration-free" SD (STFC)**
  - Useful as a bootstrap mechanism
    - it can locate the information server on the local subnet
  - Will use an existing protocol
  - May not happen – insufficient time
- **Make use of the SD APIs in all components (All!)**

INFOD-WG

- **Very R-GMA like but much more general**
- **Implementation**
  - University of Tennessee led collaboration
    - To produce and deploy an open source implementation
    - Weekly phone meetings
    - Team of 8
    - Had hoped for demo at OGF21 but did not quite make it
- **Further specification**
  - Discussed at OGF21 additions to the specification
  - Discussed “disseminator” – a bit like a Secondary Producer

Dependency Challenge and logging

- **Dependency challenge**
  - Have done the “1 week” things
  - Other items addressed by head code (with refactoring)
- **Logging**
  - Mostly conforms
- **64 Bit**
  - No problems