Data Access, Integration and Management

John Watt

http://csperkins.org/teaching/2004-2005/gc5

UNIVERSITY
of
GLASGOW

Overview

- Data Access and Integration
  - Key Requirements
- OGSA-DAI
  - Use cases
  - Architecture
  - Projects
- GT3 Data Management
  - GridFTP
  - The Reliable File Transfer Service
  - Replica Location Service

Context

- Entering an age of data
  - Data explosion
    - CERN: LHC will generate 1GB/s = 10PB/y
    - VLBA (NRAO) generates 1GB/s today
    - Pixar generate 100 TB/movie
  - Storage getting cheaper
- Data stored in many different ways
  - Data resources
    - Relational databases
    - XML databases
    - Flat files
- Need ways to facilitate
  - Data discovery
  - Data access
  - Data integration

Requirements

- **Data Sources and Resources**
  - Want access to data directly from data sources and resources
- A data **source** streams data in real (or pseudo-real) time from instruments, devices, or in-silico experiments or simulations
  - e.g. telescopes, particle colliders, remote sensors, video cameras
  - NB: Data sources may stream data for long periods of time, but some or all of the output may NOT be captured and stored in a persistent state
- A data **resource** is a persistent data stores held in file structures or a database management system (DBM)
  - Reside online in mass storage devices, or offline on magnetic media
    - Database = any organised collection of data on which operations may be performed through a defined API. Data is invariably linked (common project)
  - Want to group a logical set of data resources stored at one, or many sites
  - Must be possible to reference as a ‘virtual database’ and perform set operations (distributed data management and access)

Requirements

- **Data Structure and Representation**
  - Data is stored in a wide variety of structures, representations and technologies
  - Need to support the requirements of all scientific disciplines
  - Grid must support access to ALL types of data defined in EVERY format and representation
  - Wish to access
    - Numeric data at the highest level of precision and accuracy
    - Text data in ANT format, structure, language or coding system
    - Multimedia data in any standard (or user defined) binary format

Requirements

- **Data Organisation**
  - Traditionally, data in scientific disciplines have been organised in application specific file structures designed to optimise compute intensive data processing and analysis
    - A lot of data accessed by current Grid environments still exist in this form
  - Require the Grid to provide access to data held in DBMs and XML repositories
    - Increasingly used in bioinformatics, chemistry, environmental sciences and earth sciences
  - Why?
    - They provide the ability to store and maintain data in application-independent structures
    - They can represent data in complex ways, reflecting naturallu occuring or user-defined associations
    - Relational and object DBMs also provide facilities for automating data management and referential integrity

Requirements

- **Data Provenance**
  - A record of the origin and history of a piece of data
  - Special form of Audit trail that traces each step in sourcing, moving and processing data, together with “who did what when”
- Key to establishing ownership, quality, reliability and currency of data – particularly during discovery process
  - Also provides information that is necessary for recreating data
  - Can avoid time-consuming and resource intensive processing for recreation
- Grid must provide the capability to record data provenance, and the ability to access the provenance record
- Should be captured through automated mechanisms as far as possible
  - Should provide tools to assist owners of existing data to create important provenance elements with minimum effort

Requirements

- **Data Access Control**
  - Grid must provide controls over data access to ensure confidentiality of the data is maintained
  - Prevent users who do not have the necessary privileges to change data
- Must be possible for data owner to grant or revoke permisssions to other users (or delegate this task to a third party)
  - Common requirement for data owned or curated by an organisation
    - e.g. gene sequences, chemical structures, survey data
- Controlling the granularity of access can range from an entire database down to a subset of the data values
  - In a clinical study, for example – must be possible to access patient treatment records based on diagnosis and age range
  - View age and sex of patients without knowing their names (or doctor's name)

Requirements

- **Data Publishing and Discovery**
  - A knowledge of ownership, currency and provenance is required in order to establish the quality and reliability of the data content and so make a judgement on its value and use
  - Specification of the physical characteristics of the data e.g. volume, number of logical records and preferred access paths are necessary in order to access and transport the data efficiently
- Minimum information the user must know in order to reference a data resource is its name and location.
- A specification of its internal data structure is required in order to access its content
- Grid must provide ability to register and de-register data resources dynamically
  - Schedule these instructions, and propagate them to sites holding replicas and copies of the resources

Requirements

- **Data Operations**
- Grid must provide the ability to translate target, output and retrieval condition parameters in metadata terms into physically addressable data resources and structures
- Must construct search rules and matching criteria in the semantics and syntax of query languages form specified parameters
  - Extract data from user defined files and documents
- When more than one resource is specified, Grid must provide the ability to link them together (even if they have different data structures) to produce a single logical target that gives consistent results
  - When linking, Grid must provide ability to use data in one resource as the conditions for retrieving data from another resource (i.e. perform a sub-Query) – compare gene sequences in local DB with centralised repository

Requirements

- **Data Operations**
- Must be able to construct distributed queries when the target data resources are located at different sites, and must be able to support federated queries when some data resources are accessed through different query languages
- When metadata info is available, Grid should have the ability to override specified controls and make decisions on the preferred location and access paths to data

- Must provide ability to capture and record all observations, inferences and conclusions drawn during analysis and interpolation processes which result in data being modified

OGSA-DAI

Use Cases

- Scientist A wants to make their work available to the scientific community
  - Publish a read/only on-line database
  - Register data resource with a public registry (DATA DISCOVERY)
- Scientist B (who collaborates with A) wishes to contribute new work to A's database
  - A provides a secure private read/write grid data service (DATA ACCESS)
- Scientist C wishes to compare their work with A's but they are in different formats
  - Transform Service will translate the data into recognisable formats
  - DATA TRANSFORMATION
- Scientists X,Y and Z publish their work as read-only data resources (Z only allows specific queries to run)
  - A,B and C wish to use subsets of data from X,Y and Z….
  - DATA COLLABORATION

Use Cases

- … Scientist A writes a service which exposes an integrated set of data as another virtual data resource
  - B and C can use this resource as if it were a single data resource
    - DATA INTEGRATION

- We require middleware tools to grid-enable existing databases for discovery, access, transformation, integration and collaboration

- OGSA-DAI
  - Open Grid Services Architecture Data Access and Integration

OGSA-DAI Project

- Develop a component library
  - Access and manipulate data in a grid
  - Serve UK and International e-Science Communities
- Aims to provide
  - Common interface to data resources
  - Simple Integration of distributed queries to multiple data resources
- Contribute to standardisation efforts
  - Input into GGF DAIS WG and other groups
  - Provide a reference implementation of DAIS spec
- Based on Open Grid Services Architecture (OGSA)
  - GT3 “compliant”
- Project Partners: NeSC, EPCC, NEReSC, ESNW, IBM, Oracle
  - Funded by the Grid Core Programme - £3 million 18 months from Feb 2002
  - Three major releases, three interim releases

OGSA-DAI Project Status

- Current release 3.1
  - GT3.0.2 or 3.2 compliant
  - Platform and language independent
    - Java 1.4
    - Document model
- Work concentrated on data access
  - Wraps data resources without hiding underlying data model
  - Provides base for higher level services
    - Distributed Query Processing (DQP)
    - Data Federation services

OGSA-DAI Services

- OGSA-DAI uses three main service types
  - DAISGR (registry) for discovery
  - GDSF (factory) to represent a data source
  - GDS (data service) to access a data resource

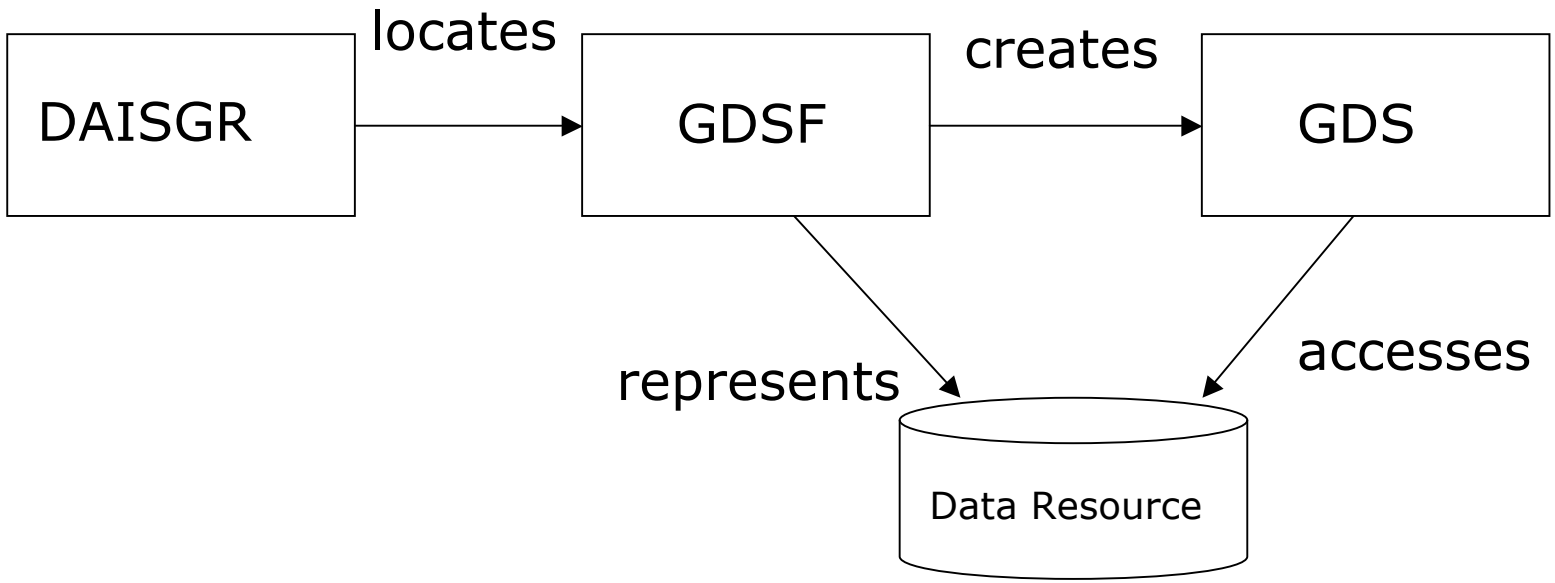

OGSA-DAI Services

- Grid Data Service Factory (GDSF)
  - Represents a data resource
  - Persistent service
    - Currently static (no dynamic GDSFs)
      - Cannot instantiate new services to represent other/new databases
  - Exposes capabilities and metadata -- May register with a DAISGR
- Grid Data Service (GDS)
  - Created by a GDSF
  - Generally transient service
  - Required to access data resource -- Holds the client session
- DAI Service Group Registry (DAISGR)
  - Persistent service, based on OGSI ServiceGroups
  - GDSFs may register with DAISGR
  - Clients access DAISGR to discover resources or services (which may need special capabilities – support a given portType or activity)

Web Services Architecture

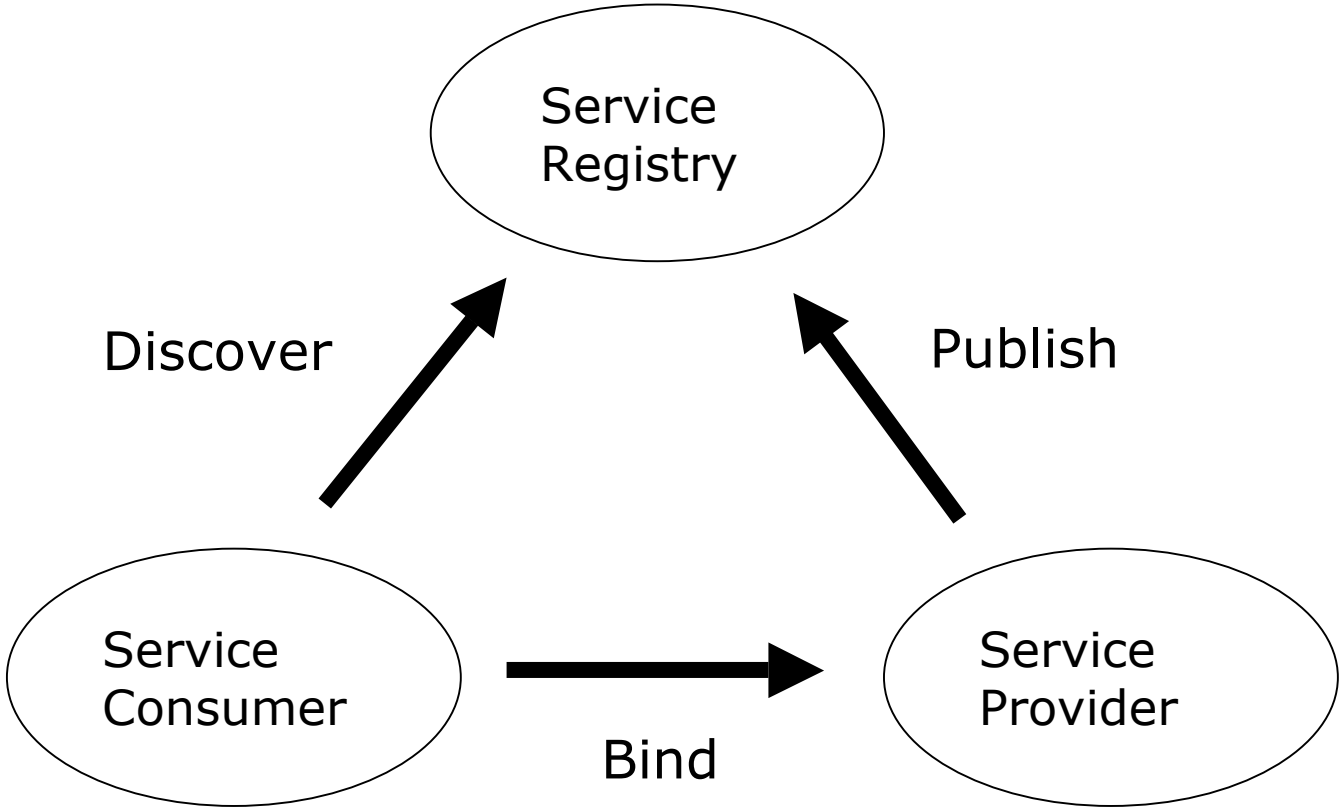

OGSA-DAI Service Architecture

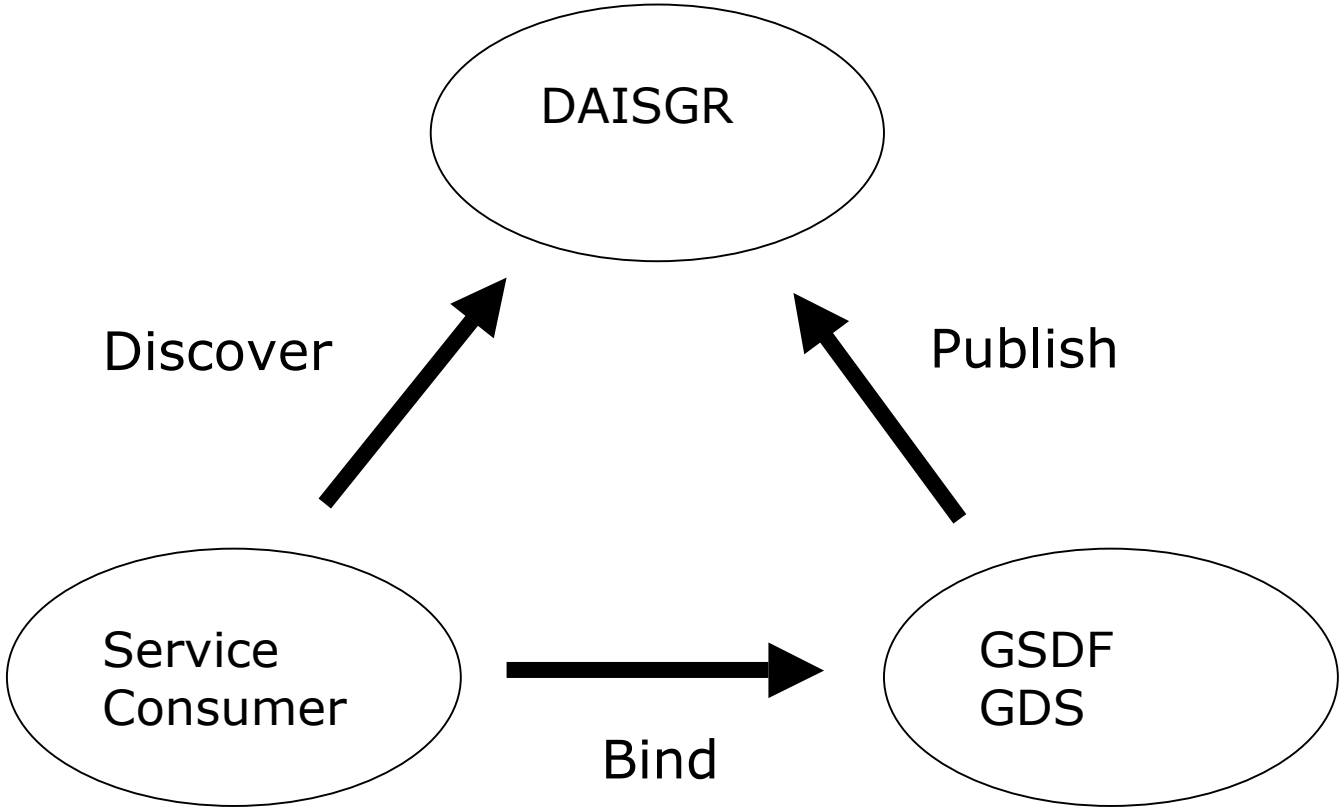

Supported Databases and Activities

- Relational
  - MySQL, DB2, Oracle
  - (postgreSQL, SQLServer)
- XML
  - Xindice
  - (eXist)
- Other
  - Files
  - ???
- Predefined Activities
  - relationalResourceManager, sqlBulkLoadRowset…
  - xmlCollectionManagement, xQueryStatement, xPathStatement
  - Deliver to URL, Deliver to GFTP …
  - outputStream, inputStream, zipArchive….

Interaction Model

- Startup
  - Start OGSI container with persistent services for DAISGR and GDSF (which represents our database)
- Registration
  - GDSF registers with DAISGR (passes GSH to DAISGR)
- Discovery
  - Client wishes to access database, either
    - Query the GDSF directly if known, or
    - Identify suitable GDSF through DAISGR (DAISGR passes GSH to client)
- Service Creation
  - Having identified a suitable GDSF, client asks the GDSF to create a GDS (createService)

Interaction Model

- Perform
  - Client interacts with the GDS by sending Perform documents
  - The GDS responds with a Response document
  - Client may terminate the GDS when finished or let it die naturally

- Summary
  - Only describe an access use case
    - Client not concerned with connection mechanism
    - Similar framework could accommodate service-service interactions
  - Discovery aspect is important
    - Probably requires a human
    - Needs adequate definition of metadata
      - Definitions of ontologies and vocabularies – not something OGSA-DAI does

OGSA-DAI Projects

- ODD-Genes
  - Data Analysis for genetics
  - http://www.epcc.ed.ac.uk/oddgenes
- FirstDIG
  - Data mining with the First Transport Group, UK
    - Example: “When buses are more than 10 minutes late there is an 82% chance that revenue drops by at least 10%”
    - OGSA-DAI used to access data on Customers, Mileage, Schedule Adherance and revenue. Data mining app accessesthis through OGSA-DAI client app
  - http://www.epcc.ed.ac.uk/firstdig
- Other projects
  - MCS on OGSA-DAI, BioGrid, OpenSkyQuery
  - http://www.ogsadai.org.uk/projects/

Globus Toolkit Data Management

- The third and final pillar in Globus Toolkit

Data Management

Globus Toolkit Data Management

- Components of GT3 Data Management:
  - GridFTP
    - A high performance, secure, reliable data transfer protocol optimised for high-bandwidth, wide-area networks.
  - Reliable File Transfer
    - OGSA based service providing interfaces for controlling and monitoring 3rd party file transfers using GridFTP servers
  - Replica Location Service
    - Maintains and provides access to mapping information from logical names for data items to target names
  - XIO
    - An extensible input/output library for the Globus Toolkit that provides a simple and intuitive API (open/close/read/write) to swappable IO implementations

GridFTP

- GridFTP protocol is based on FTP
  - Have selected a set of protocol features and extensions already defined in the IETF RFCs and added a few additional grid-like features

- GridFTP provides the following
  - GSI security on control and data channels
  - Multiple data channels for parallel transfers
  - Partial file transfers
  - Third-party (direct server-server) transfers
  - Authenticated data channels
  - Reuseable data channels
  - Command pipelining

Reliable File Transfer

- OGSI compliant service
- Uses existing GridFTP (non-OGSI) protocols and tools to execute 3rd party transfer for the user
- Provides extensive state transition notification

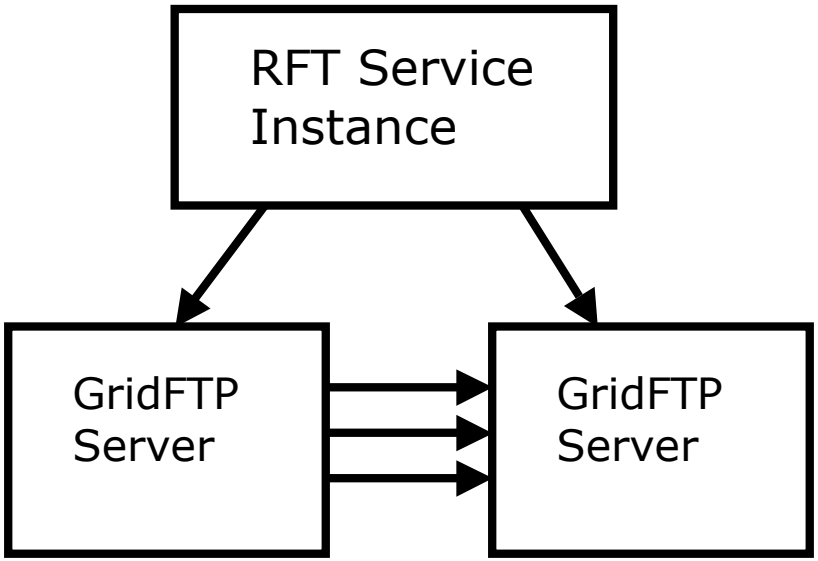

RFT on the move…

- A Grid service container is started up, containing an RFT factory
- The RFT Factory registers itself (1) with a known registry
- A client discovers an RFT factory by querying the serviceData of the registry (2)
- Client calls the createService operation on the factory and passes in a transfer request (3)

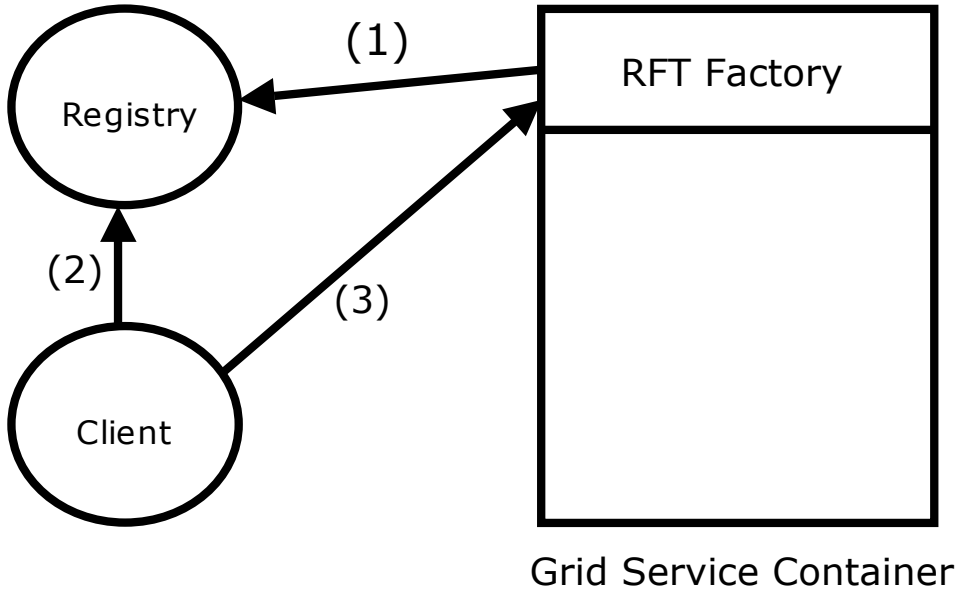

RFT on the move…

- The service instance is started (4)
- The RFT factory returns a locator (5) to the client
- Client calls Start() (6), and subscribes to notifications etc.

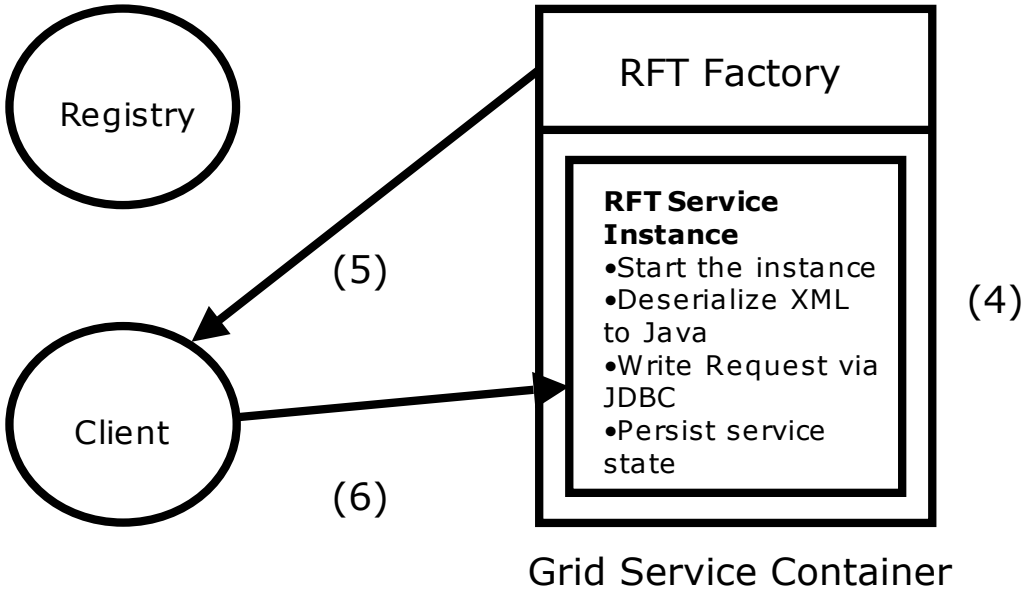

RFT Design Issues

- Lifetime management
  - key aspect of OGSI
  - Was not intuitively clear how to handle this for “disconnected” services
    - The Globus Alliance (perhaps not optimal) solution is to give it an indefinite lifetime
  - Should there be an activity monitor?
    - Does that really solve the problem? Any other ideas?

- Data virtualization
  - We virtualize data movement
    - LBL and ANL have working interoperable implementation.
    - Need to standardize…
  - Data Services Virtualization from DAIS
    - Should we pass around GSHs of file services rather than URLs?
  - Granularity
    - Single file .vs. Multi-file .vs. service composition

RFT Design Issues

- Use Service Data Elements (SDEs)
- A huge improvement over the non-OGSI services
- Information Services are (should be) baked in to the services
- Defined both push (event notification) and pull (full transfer status) SDEs
- Need to be cognizant of size, frequency, and performance of notifications
- Lots of interesting possibilities: bandwidth, errors, network status, etc..

- Design still ongoing…

Replica Location Service

- Replication of data items can
  - Reduce access latency
  - Improve data locality
  - Increase robustness, scalability and performance for distributed applications
- Replica Location Service (RLS) co-developed by the Globus team and WP2 of the DataGrid project
- The RLS typically functions as one component of a data grid
  - Along with RFT, metadata management, reliable replication and workflow management