Enabling Grids for E-science in Europe

www.eu-egee.org

Grid Summer School
Vico Equense, 27 July 2004

An Introduction to Data Services & OGSA-DAI

Malcolm Atkinson
Director
National e-Science Centre
www.nesc.ac.uk

EGEE is a project funded by the European Union under contract IST-2003-508833

Outline

Enabling Grids for E-science in Europe

- **Outline of OGSA-DAI day**
- **What is e-Science?**
  - **Collaboration & Virtual Organisations**
  - **Structured Data at its Foundation**
- **Motivation for DAI**
  - **Key Uses of Distributed Data Resources**
  - **Challenges**
  - **Requirements**
- **Standards and Architectures**
  - **OGSA Working Group**
  - **DAIS Working Group**
- **Introduction to DAI**
  - **Conceptual Models**
  - **Architectures**
  - **Current OGSA-DAI components**

Workshop Overview

OGSA-DAI Workshop

National e-Science Centre

09:00 Introduction: Data Access & Integration
Malcolm Atkinson

10:30 OGSA-DAI Tutorial: Grid Data Service
Tom Sudgen

11:00 – Coffee break

11.30 OGSA-DAI Tutorial Continues: Tom Sudgen

OGSA-DAI Users Guide:
Client-toolkit APIs
OGSA-DAI support and examples

13:00 LUNCH

15:00 Practical Part 1: Data Browser & Client Toolkit
Tom Sugden assisted by: Guy Warner & Malcolm Atkinson

17:00 BREAK

17:30 Practical Part 2: Advanced use of Client Toolkit
Tom Sugden assisted by: Guy Warner & Malcolm Atkinson

19:00 End of Lab sessions

epcc

IBM

e-Science dti

The OGSA-DAI Team

National e-Science Centre

Mario Antonioletti	EPCC	Malcolm Atkinson	NeSC
Rob Baxter	EPCC	Andrew Borley	IBM
Neil Chue Hong	EPCC	Brian Collins	IBM
Jonathan Davies	IBM	Ally Drumshushoa	EPCC
Desmond Fitzgerald	Manchester Uni.	Alastair Hume	IBM
Mike Jackson	EPCC		
Kostas Karassavas	NeSC	Amrey Krause	EPCC
Andy Laws	IBM	Charaka P	EPCC
Norman Paton	Manchester Uni.	Dave Pearson	Oracle
Tom Sudgen	EPCC	Dave V	
Dave Watson	IBM	Paul Watson	Newcastle University

e-Science & Data

What is e-Science?

- **Goal: to enable better research**
- **Method: Invention and exploitation of advanced computational methods**
  - to generate, curate and analyse research data
    - From experiments, observations and simulations
    - Quality management, preservation and reliable evidence
  - to develop and explore models and simulations
    - Computation and data at extreme scales
    - Trustworthy, economic, timely and relevant results
  - to enable *dynamic* distributed virtual organisations
    - Facilitating collaboration with information and resource sharing
    - Security, reliability, accountability, manageability and *agility*

Multiple, independently managed sources of data – each with own time-varying structure

Creative researchers discover new knowledge by combining data from multiple sources

The Primary Requirement …

Enabling Grids for E-science in Europe

Enabling *People* to Work Together on Challenging Projects: Science, Engineering & Medicine

Three-way Alliance

Multi-national, Multi-discipline, Computer-enabled Consortia, Cultures & Societies

Theory
Models & Simulations
→
Shared Data

Experiment &
Advanced Data
Collection
→
Shared Data

Computing Science
Systems, Notations &
Formal Foundation
→ Process & Trust

Requires Much Engineering, Much Innovation

Changes Culture, New Mores, New Behaviours

New Opportunities, New Results, New Rewards

Examples of DAI

Composing Observations in Astronomy

Enabling Grids for E-science in Europe

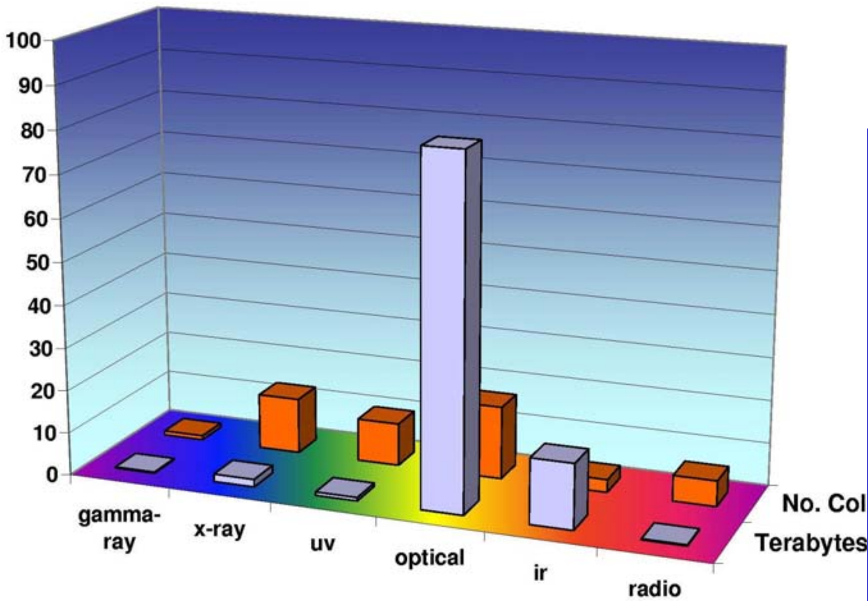

No. & sizes of data sets as of mid-2002, grouped by wavelength

- 12 waveband coverage of large areas of the sky
- Total about 200 TB data
- Doubling every 12 months
- Largest catalogues near 1B objects

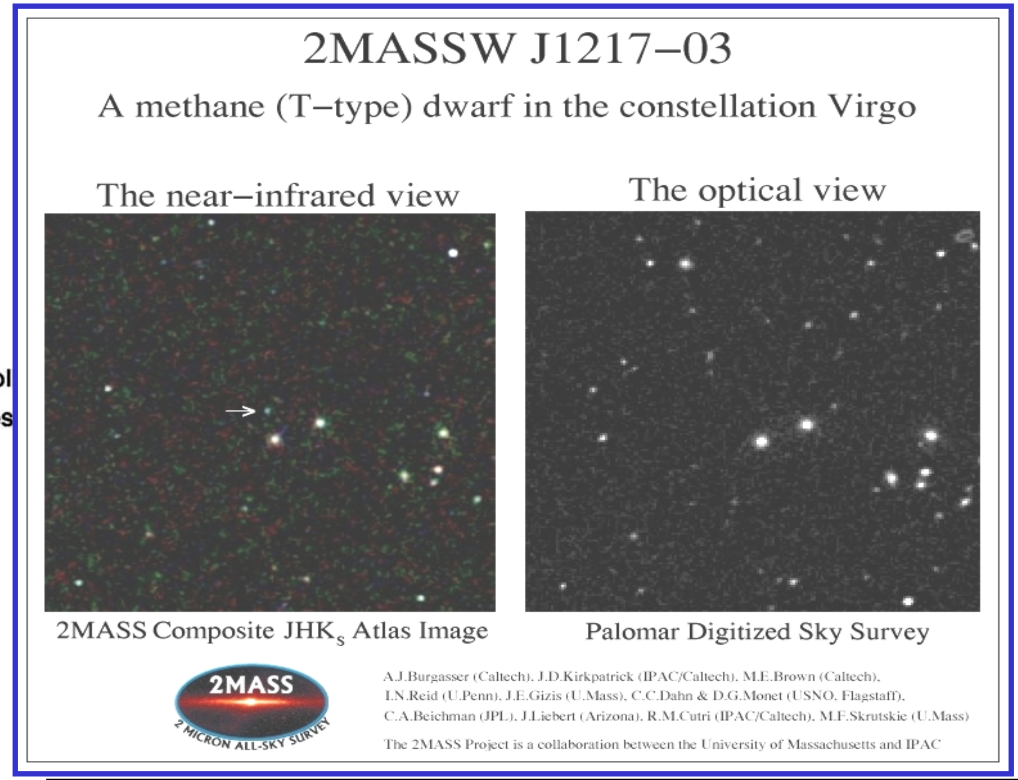

THE NEW YORK TIMES NATIONAL TUESDAY, JUNE 1, 1999

Astronomers Detect New Category of Elusive 'Brown Dwarfs'

By JOHN NOBLE WILFORD

CHICAGO, May 31 — Ambitious

Apache Point, N.M. Dr. Michael Strauss and a graduate student,

was a brown dwarf, but was not associated with a star companion.

have been possible in hotter, younger objects. An estimate of their mass

Data and images courtesy Alex Szalay, John Hopkins

Wellcome Trust: Cardiovascular Functional Genomics

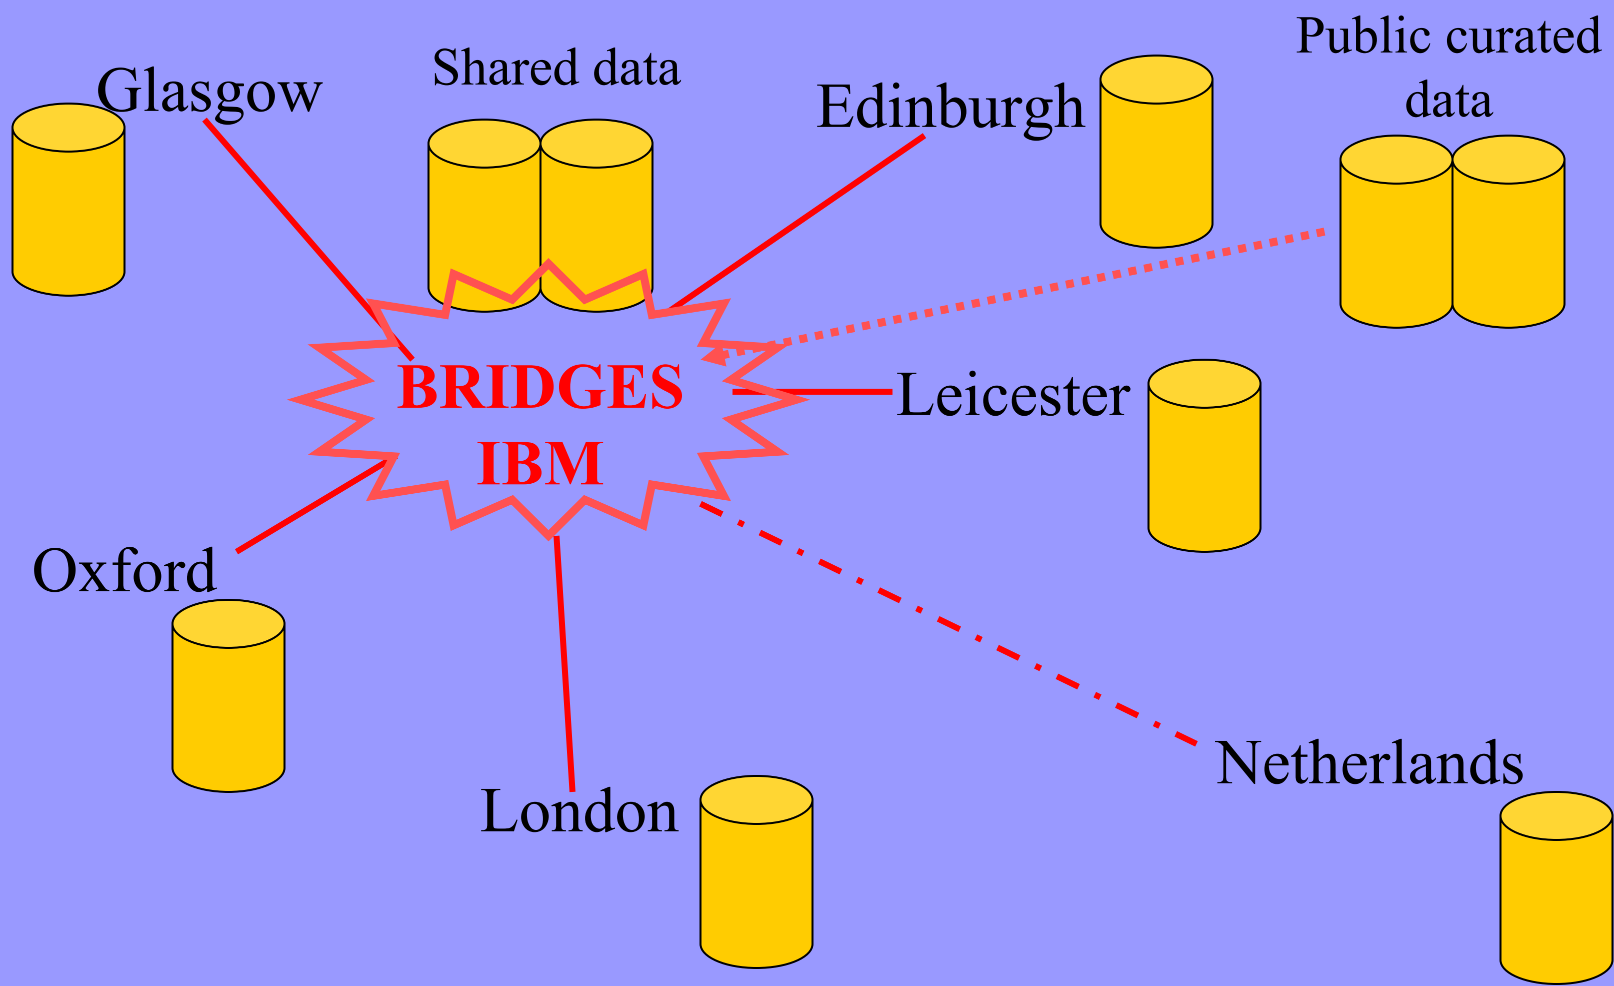

the globus project™
www.globus.org

The Emergence of Global Knowledge Communities

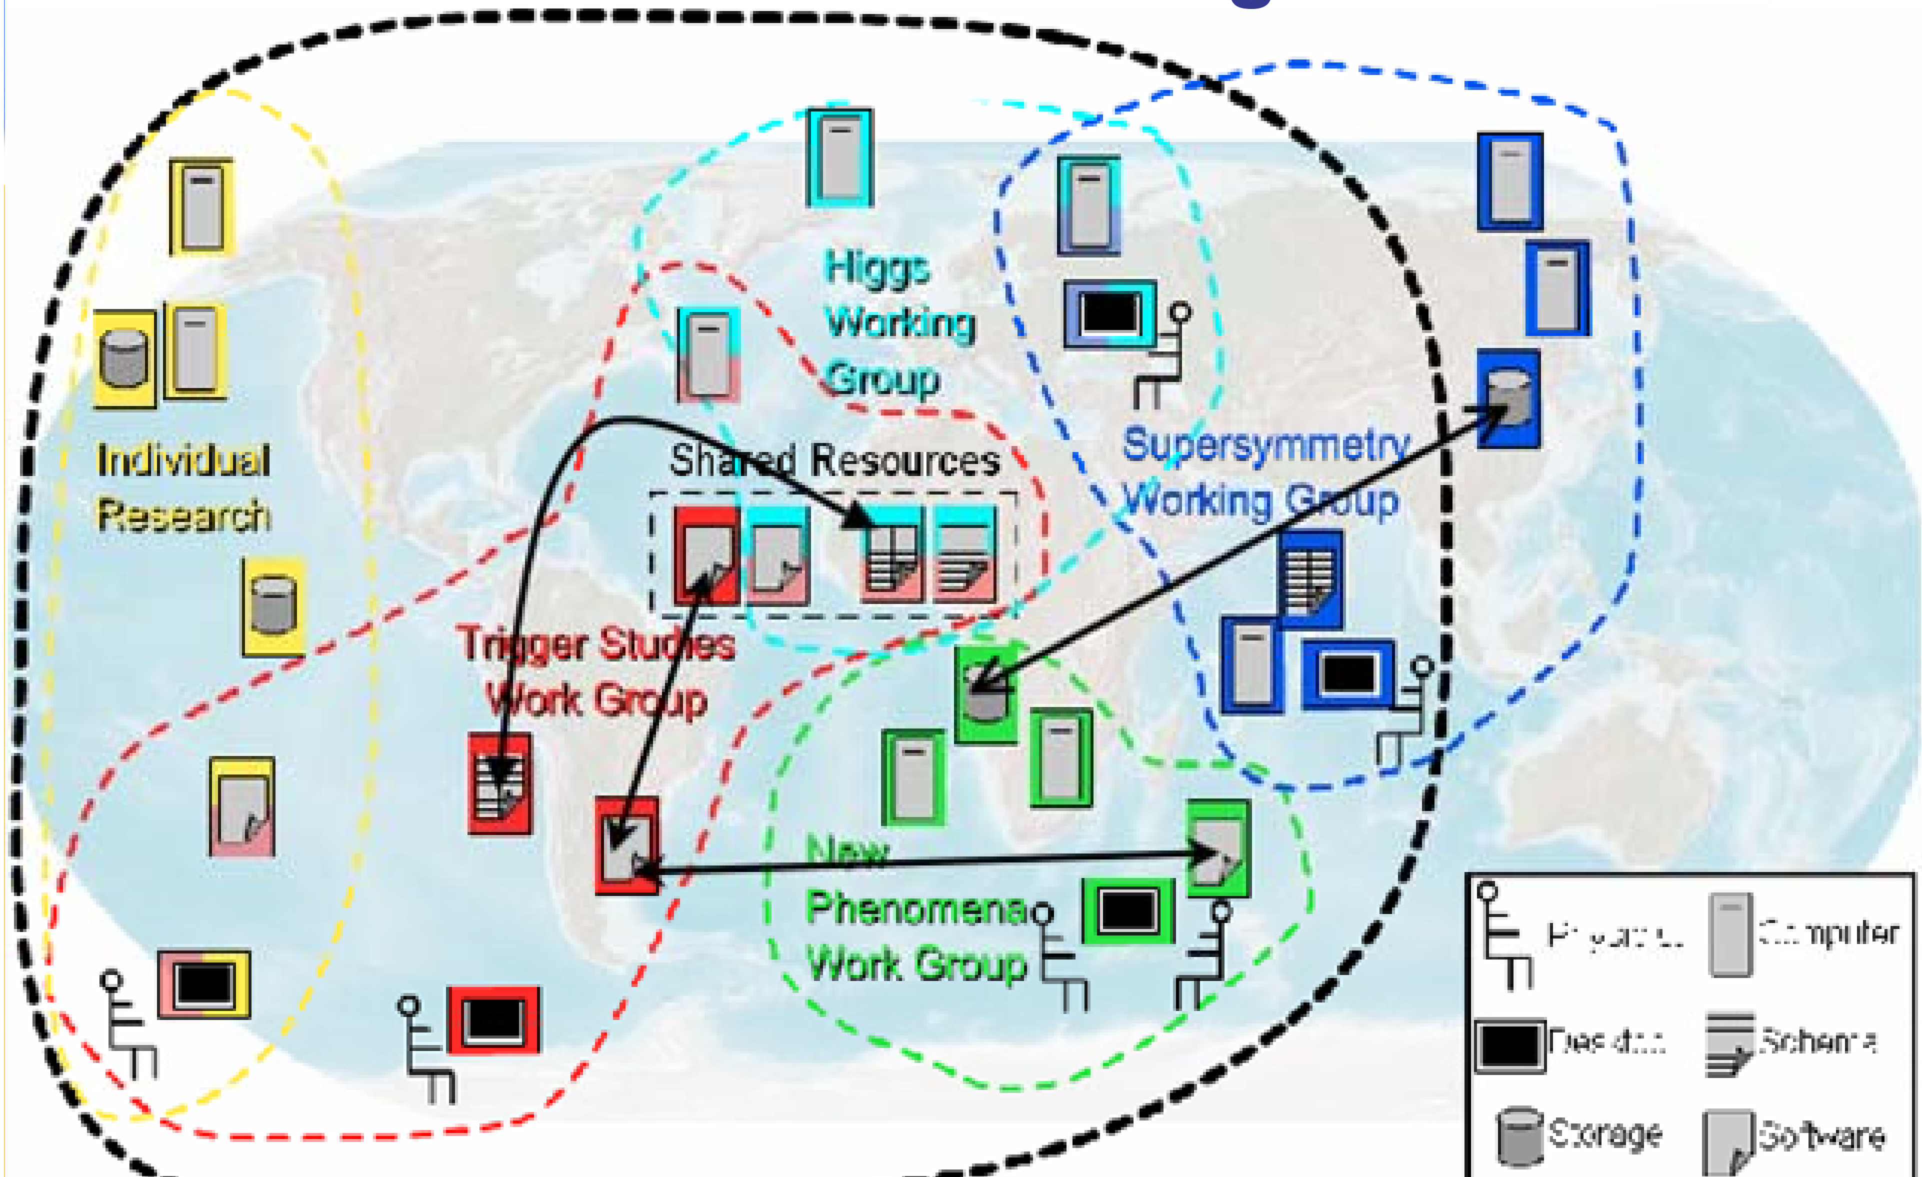

Slide from Ian Foster's ssdbm 03 keynote

Database Growth

Bases 41,073,690,490

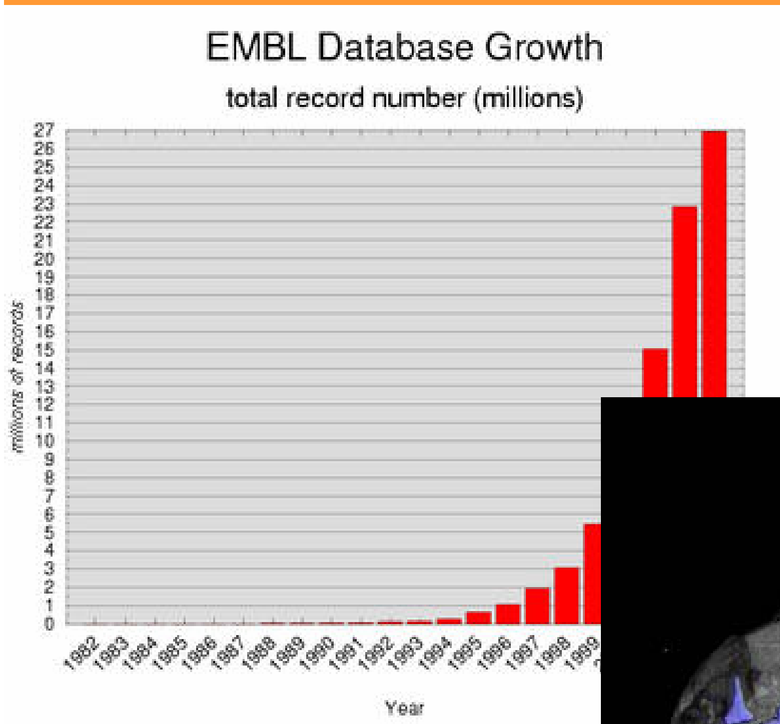

PDB Content Growth

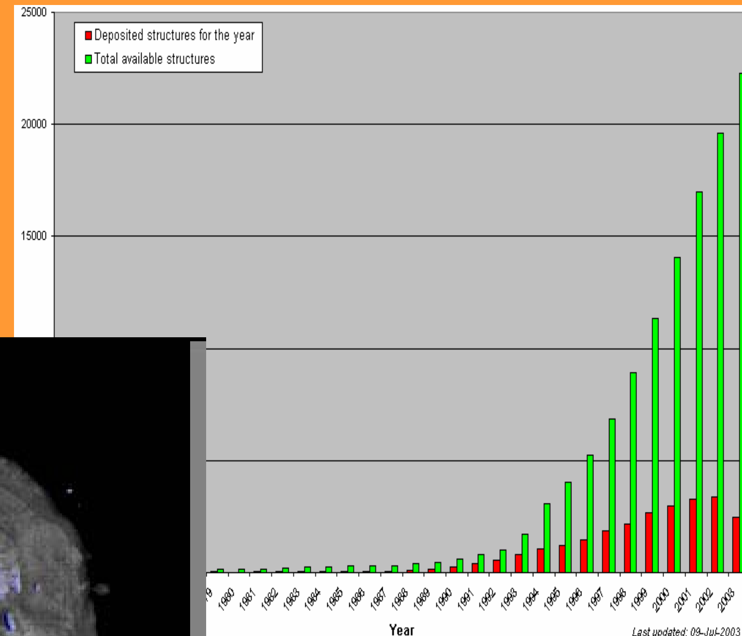

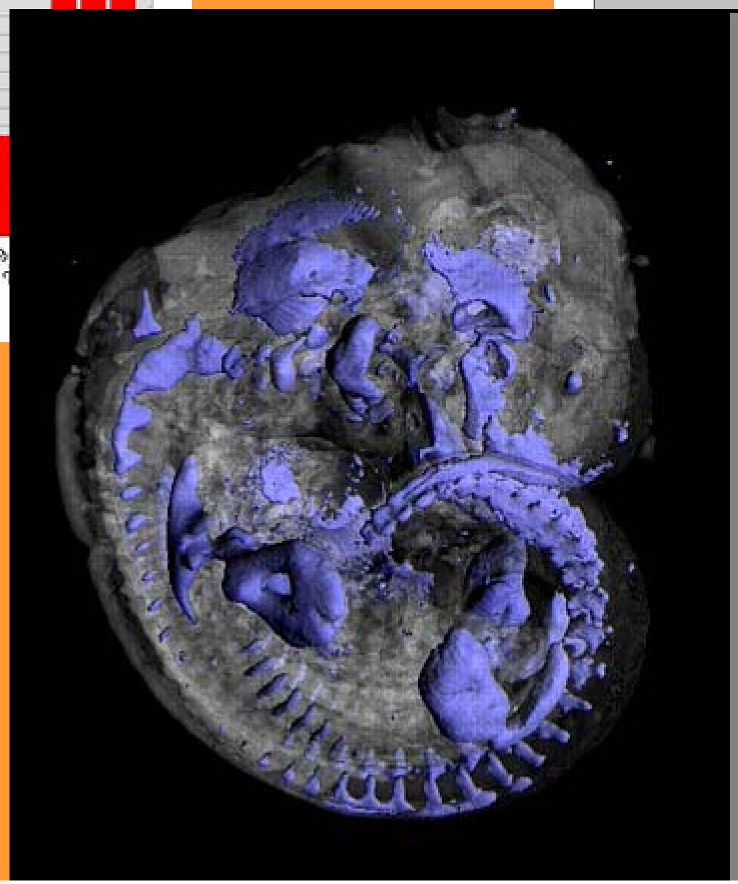

MRC Medical Research Council

MRC HUMAN GENETICS UNIT · EDINBURGH ·

epcc

Biochemical Pathway Simulator

(Computing Science, Bioinformatics, Beatson Cancer Research Labs)

Walter Kolch

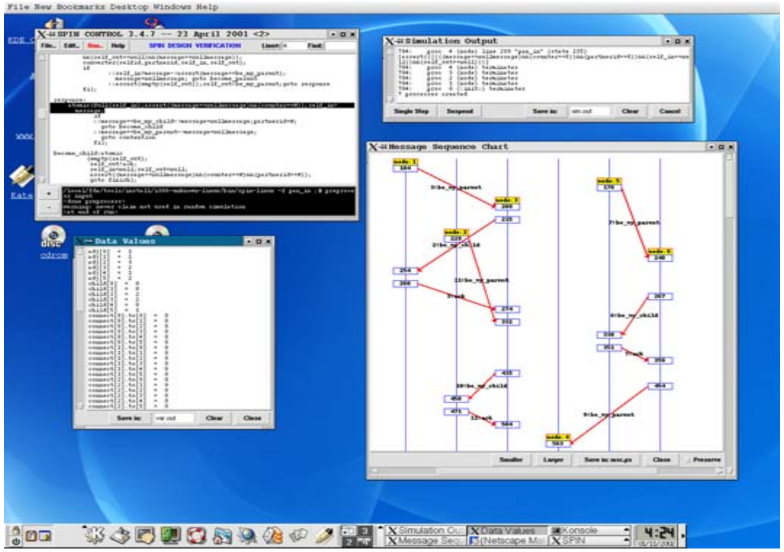

Closing the information loop – between lab and computational model.

DTI Bioscience *Beacon Project*

Now largest EU project in the Life Sciences – see
http://www.cancerresearchuk.org/news/pressreleases/scottishscientists_22july04

DAI: What is needed?

Data Access and Integration: motives

- **Key to Integration of Scientific Methods**
  - Publication and sharing of results
    - Primary data from observation, simulation & experiment
    - Encourages novel uses
    - Allows validation of methods and derivatives
    - Enables discovery by combining data independently collected
- **Key to Large-scale Collaboration and Decisions!**
  - **Economies: data production, publication & management**
    - Sharing cost of storage, management and curation
    - Many researchers contributing increments of data
    - Pooling annotation ≡ rapid incremental publication
    - And criticism
  - **Accommodates global distribution**
    - Data & code travel faster and more cheaply
  - **Accommodates temporal distribution**
    - Researchers assemble data
    - Later (other) researchers access data

Responsibility
Ownership
Credit
Citation
?

Data Access and Integration: challenges

Petabyte of Digital Data / Hospital / Year

- **Scale**
  - **Many sites, large collections, many uses**
- **Longevity**
  - **Research requirements outlive technical decisions**
- **Diversity**
  - **No "one size fits all" solutions will work**
  - **Primary Data, Data Products, Meta Data, Administrative data, …**
- **Many Data Resources**
  - **Independently owned & managed**
    - **No common goals**
    - **No common design**
    - **Work hard for agreements on foundation types and ontologies**
    - **Autonomous decisions change data, structure, policy, …**
  - **Geographically distributed**

Data Access and Integration: Scientific discovery

Enabling Grids for E-science in Europe

- **Choosing data sources**
  - **How do *you* find them?**
  - **How do *they* describe and advertise them?**
  - **Is the equivalent of Google possible?**
- **Obtaining access to that data**
  - **Overcoming administrative barriers**
  - **Overcoming technical barriers**
- **Understanding that data**
  - **The parts *you* care about for *your* research**
- **Extracting nuggets from multiple sources**
  - **Pieces of *your* jigsaw puzzle**
- **Combing them using sophisticated models**
  - **The *picture* of reality in *your* head**
- **Analysis on scales required by statistics**
  - **Coupling data access with computation**
- **Repeated Processes**
  - **Examining variations, covering a set of candidates**
  - **Monitoring the emerging details**
  - **Coupling with scientific workflows**

You're an innovator

$\therefore$ Your model $\neq$ their model

$\Rightarrow$ Negotiation & patience needed from *both* sides

Tera → Peta Bytes

Enabling Grids for E-science in Europe

- RAM time to move
  - 15 minutes
- 1Gb WAN move time
  - 10 hours ($1000)
- Disk Cost
  - 7 di… = $5…
- Disk P…
  - 10…
- Disk W…
  - 5.6 Kg
- Disk Footprint
  - Inside machine

- RAM time to move
  - 2 months
- 1Gb WAN move time
  - 14 month…
- … C…st
  - …0 D…k… …0 units +
  - … r…cks = $7 million
- Disk Power
  - 100 Kilowatts
- Disk Weight
  - 33 Tonnes
- Disk Footprint
  - 60 m²

Now make it secure & reliable!

May 2003 Approximately Correct *Distributed Computing Economics*

Jim Gray, Microsoft Research, MSR-TR-2003-24

Mohammed & Mountains

- **Petabytes of Data cannot be moved**
  - It stays where it is produced or curated
    - Hospitals, observatories, European Bioinformatics Institute, …
  - A *few* caches and a *small* proportion cached
- **Distributed collaborating communities**
  - Expertise in curation, simulation & analysis
- **Distributed & diverse data collections**
  - Discovery depends on insights
    - ⇒ Unpredictable sophisticated application code
  - Tested by combining data from many sources
  - Using *novel* sophisticated models & algorithms
- **What can you do?**

Architectural Requirement: Dynamically
Move computation to the data

- **Assumption: code size << data size**
- **Develop the database philosophy for this?**
  - **Queries are dynamically re-organised & bound**
- **Develop the storage architecture for this?**
  - **Compute closer to disk?**
    - **System on a Chip using free space in the on-disk controller**
- **Data Cutter a step in this direction**
- **Develop experiment, sensor & simulation architectures**
  - **That take code to select and digest data as an output control**
- **Safe hosting of arbitrary computation**
  - ***Proof-carrying code* for data and compute intensive tasks + *robust hosting* environments**
- **Provision combined storage & compute resources**
- **Decomposition of applications**
  - **To ship behaviour-bounded sub-computations to data**
- **Co-scheduling & co-optimisation**
  - **Data & Code (movement), Code execution**
  - **Recovery and compensation**

Dave Patterson
Seattle
SIGMOD 98

Little is done yet – requires much R&D and a Grid infrastructure

Scientific Data: Opportunities and Challenges

Enabling Grids for E-science in Europe

- **Opportunities**
  - Global Production of *Published* Data
  - Volume↑ Diversity↑
  - Combination ⇒ Analysis ⇒ Discovery
- **Challenges**
  - Data Huggers
  - Meagre metadata
  - Ease of Use
  - ...tegration
  - Depen...

A Cornucopia of Research Challenges

- **Opportunities**
  - Sp... ...xing
  - New Data Organisation
  - New Algorithms
  - Varied Replication
  - Shared Annotation
  - Intensive Data & Computation
- **Challenges**
  - Fundamental Principles
  - Approximate Matching
  - Multi-scale optimisation
  - Autonomous Change
  - Legacy structures
  - Scale and Longevity
  - Privacy and Mobility
  - Sustained Support / Funding

The Story so Far

- **Technology enables Grids, More Data & …**
- **Distributed systems for sharing information**
  - **Essential, ubiquitous & challenging**
  - **Therefore share methods and technology as much as possible**
- **Collaboration is essential**
  - **Combining approaches**
  - **Combining skills**
  - **Sharing resources**
- **(Structured) Data is the language of Collaboration**
  - **Data Access & Integration a Ubiquitous Requirement**
  - **Primary data, metadata, administrative & system data**
- **Many hard technical challenges**
  - **Scale, heterogeneity, distribution, dynamic variation**
- **Intimate combinations of data and computation**
  - **With unpredictable (autonomous) development of both**

Structure enables understanding, operations, management and interpretation

Outline

Enabling Grids for E-science in Europe

- Outline of OGSA-DAI day
- What is e-Science?
  - Collaboration & Virtual Organisations
  - Structured Data at its Foundation
- Motivation for DAI
  - Key Uses of Distributed Data Resources
  - Challenges
  - Requirements
- **Standards and Architectures**
  - **OGSA Working Group**
  - **DAIS Working Group**
- **Introduction to DAI**
  - **Conceptual Models**
  - **Architectures**
  - **Current OGSA-DAI components**

GGF OGSA WG

OGSA WG overview & goals

- **Open Grid Services Architecture**
  - **NOT** Open Grid Services Infrastructure (OGSI)
  - Seeking an Integrated Framework
  - For all Grid Functionality
- **Goal: A high-level description**
  - Functionality of components / protocols
  - Standard patterns
  - Minimum required behaviour
- **Partitioned Functions**
  - Execution Management Services
  - Data Services
  - Resource Management Services
  - Security Services
  - Self-Management Services
  - Information Services

Three useful documents:
Use cases
Glossary
Architecture
(draft-ggf-ogsa-spec-019)
http://forge.gridforum.org/projects/ogsa-wg

Scope of OGSA

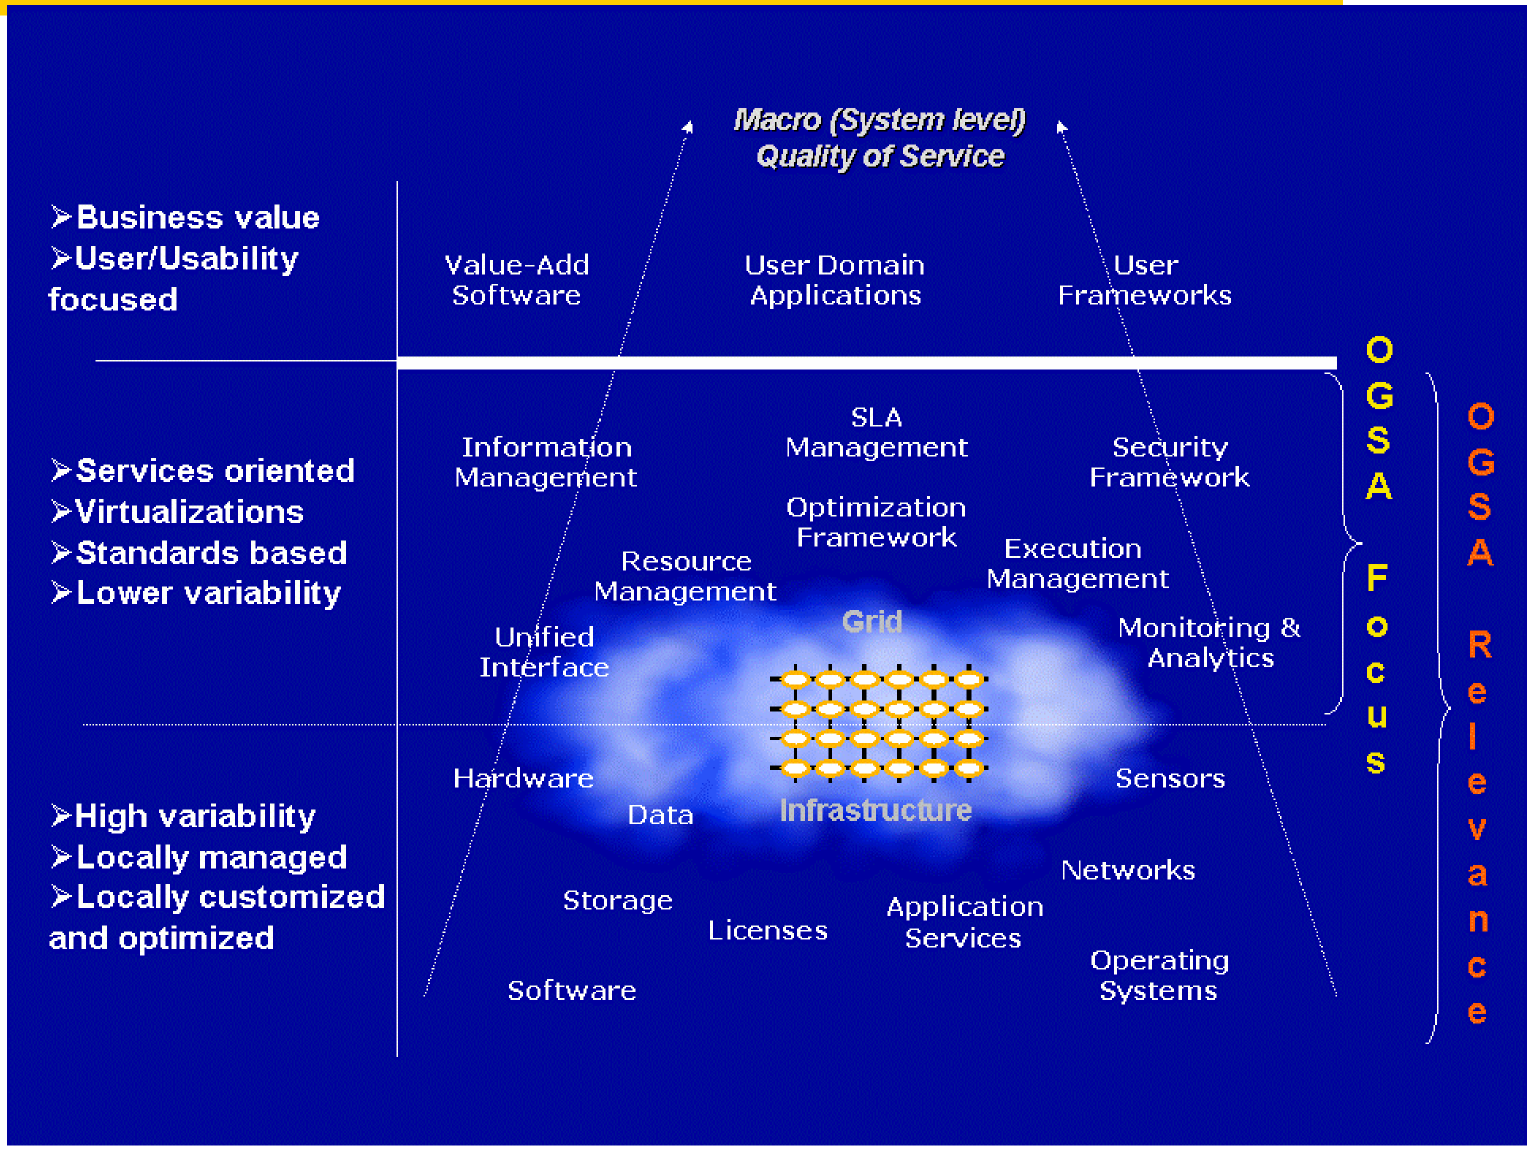

Partitioning the Scope of OGSA

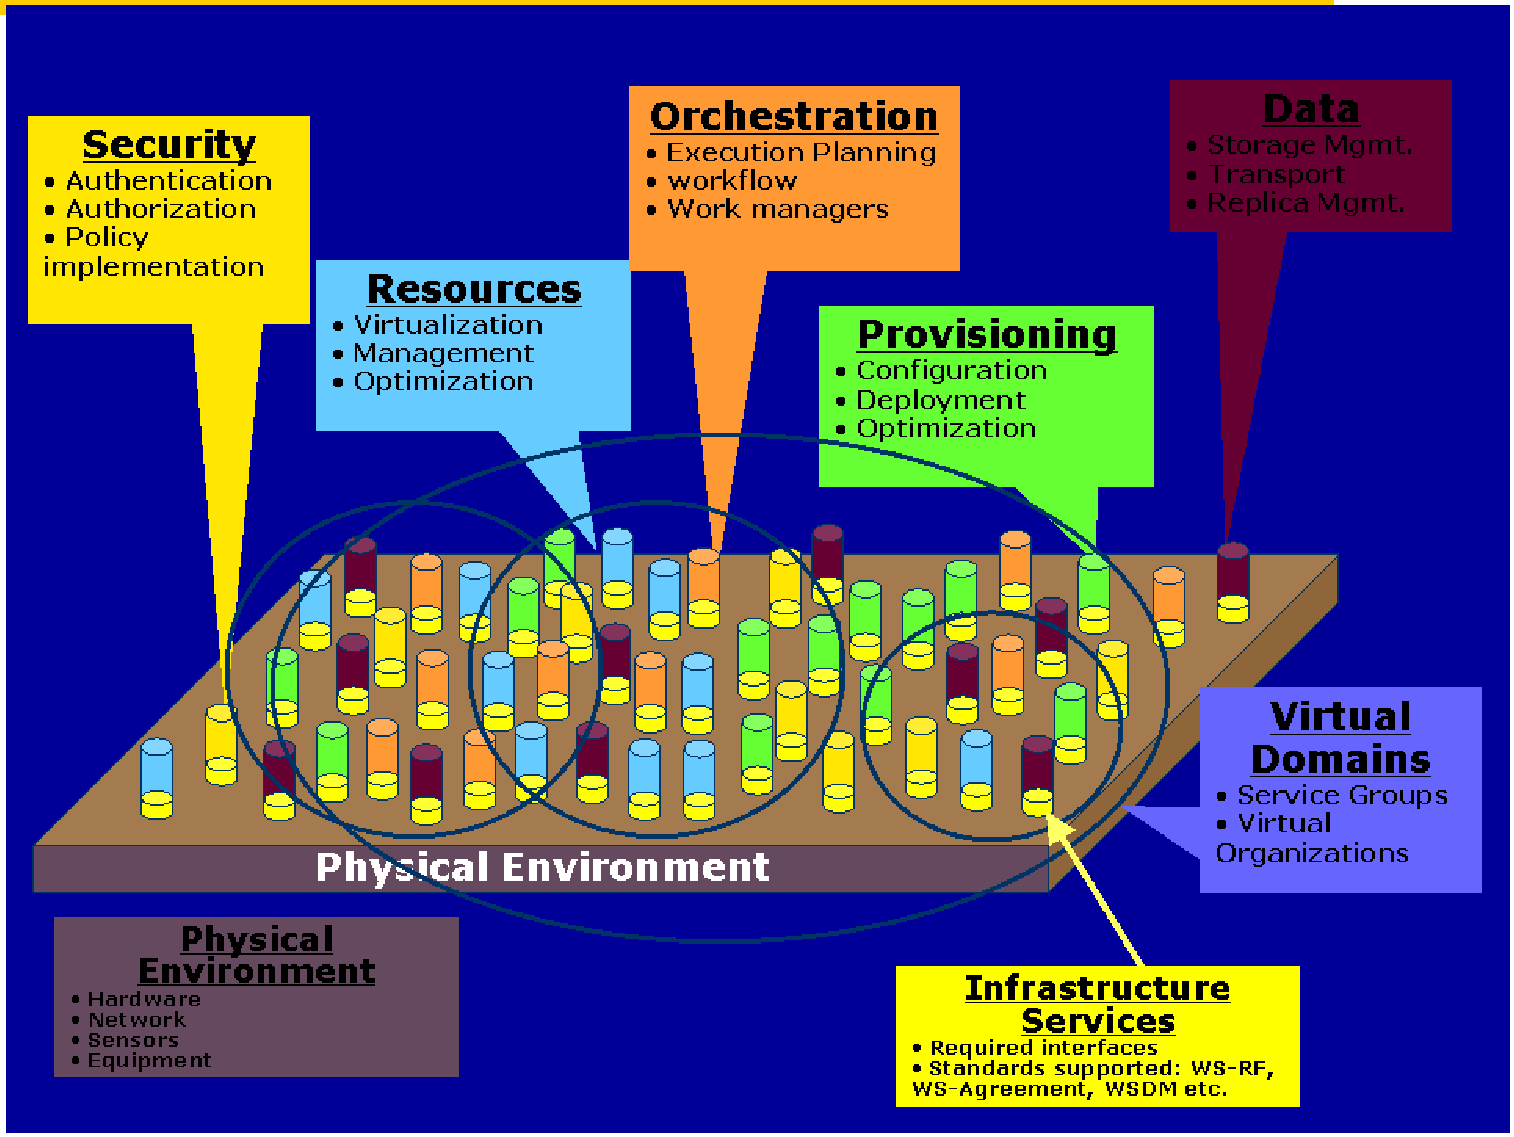

OGSA Data Services Patterns

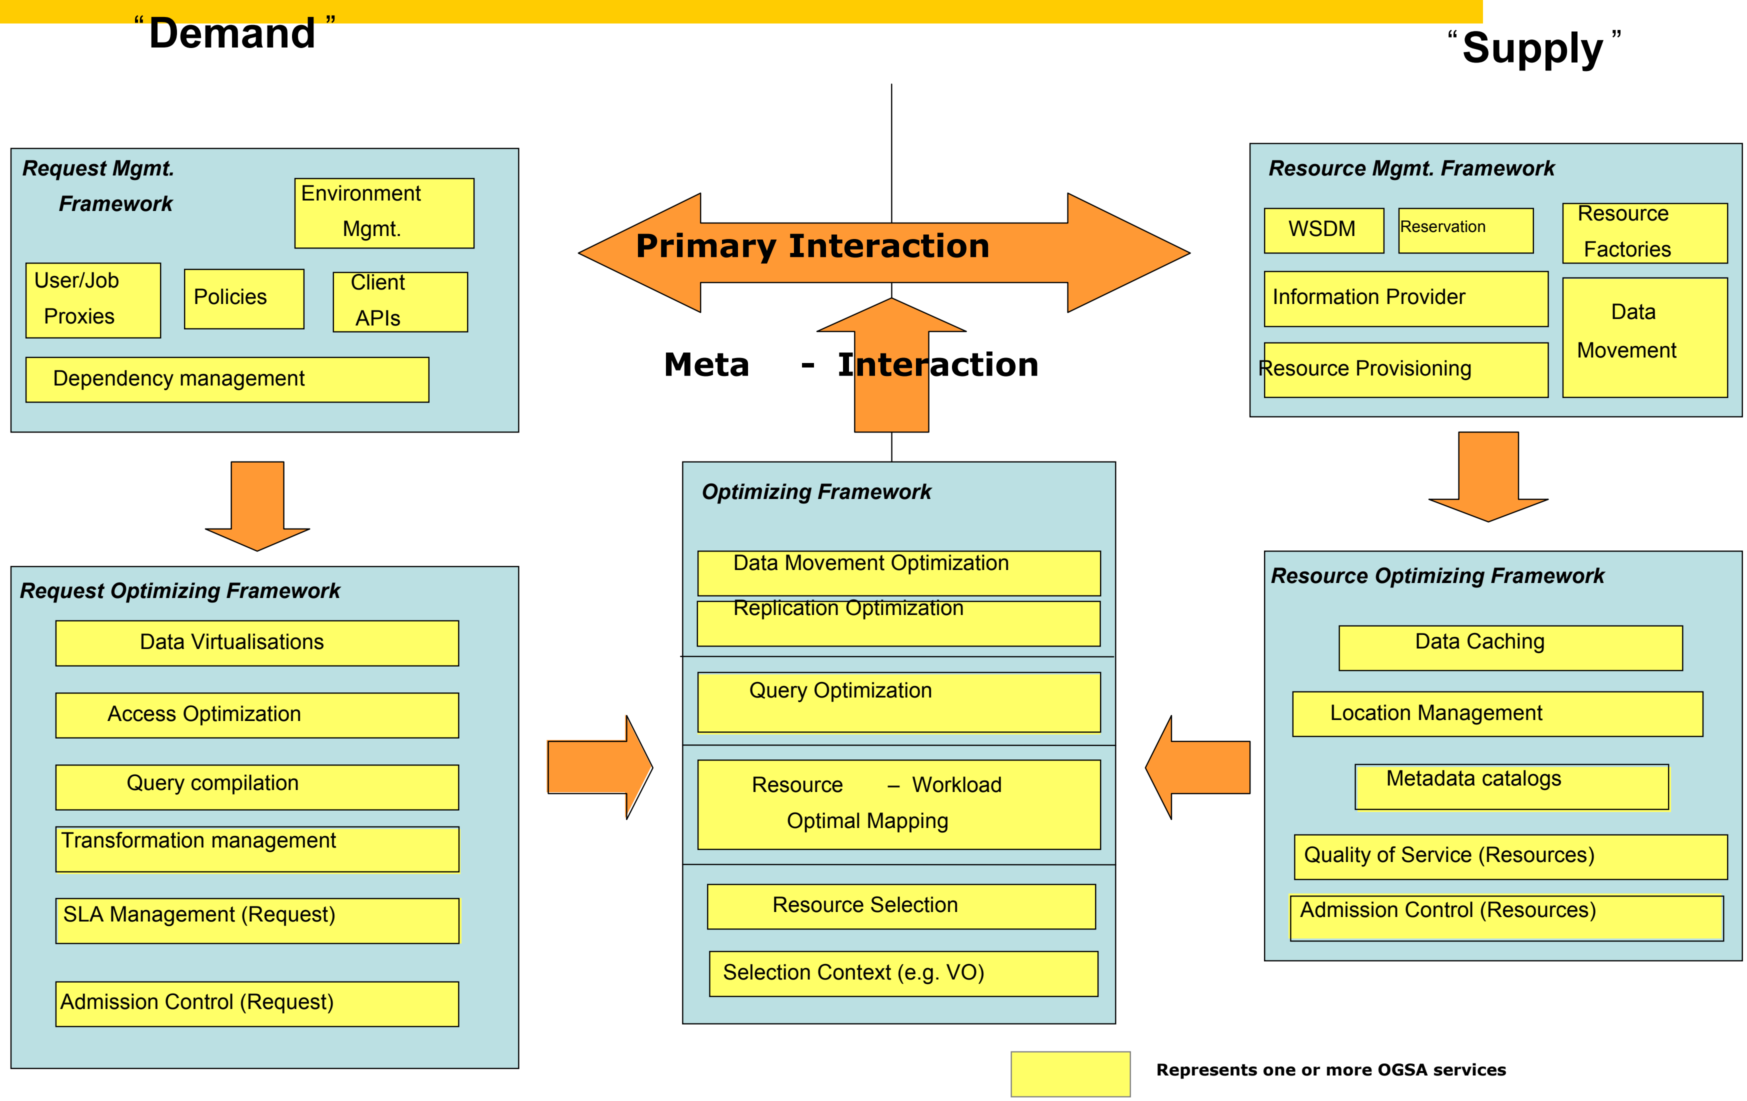

DAIS WG

DAIS WG Goals

- Provide service-based access to structured data resources as part of OGSA architecture
- Specify a selection of interfaces tailored to various styles of data access starting with relational and XML
- Interact well with other GGF OGSA specs

DAIS WG Non-Goals

- No new common query language
- No schema integration or common data model
- No common namespace or naming scheme
- No data resource management
  - E.g starting/stopping database managers
- No push based delivery
  - Information Dissemination WG?

That doesn’t mean you wont need them!
http://forge.gridforum.org/projects/dais-wg

DAIS View Of Data Services Model

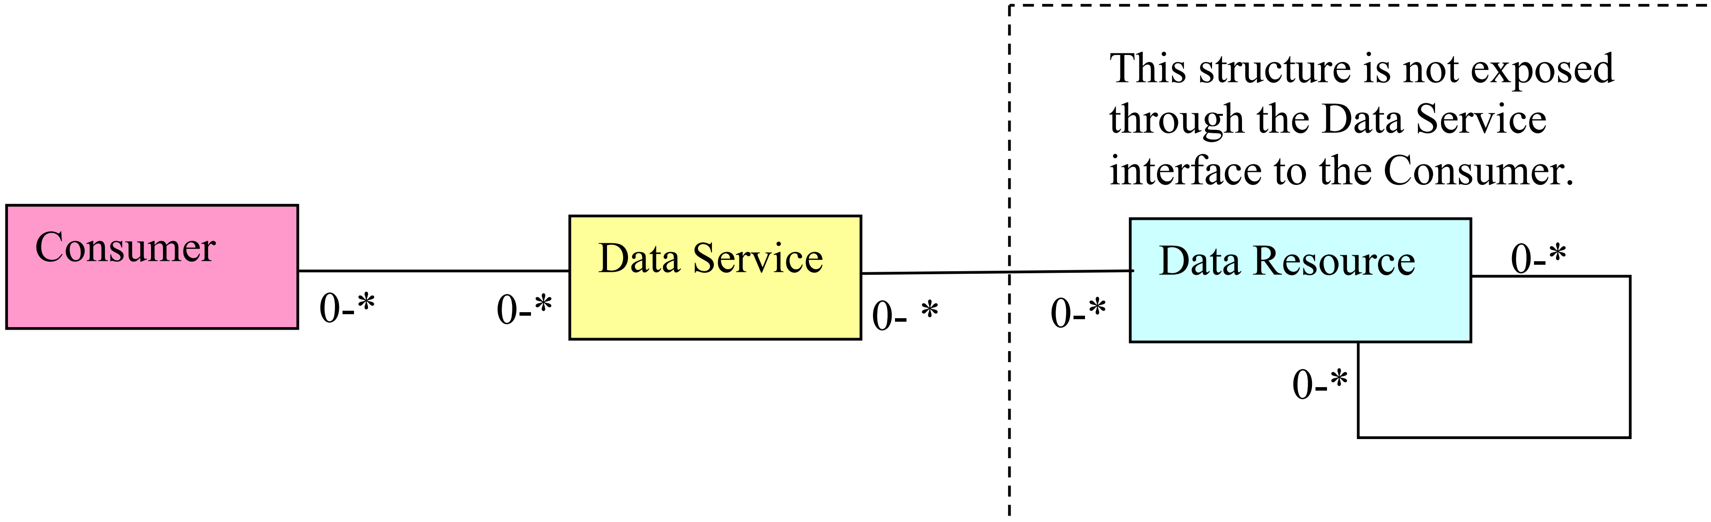

A Data Service presents a Consumer with an interface to a Data Resource. A Data Resource can have arbitrary complexity, for example, a file on an NFS mounted file system or a federation of relational databases. A Consumer is not typically exposed to this complexity and operates within the bounds and semantics of the interface provided by the Data Service

Specifying Interfaces

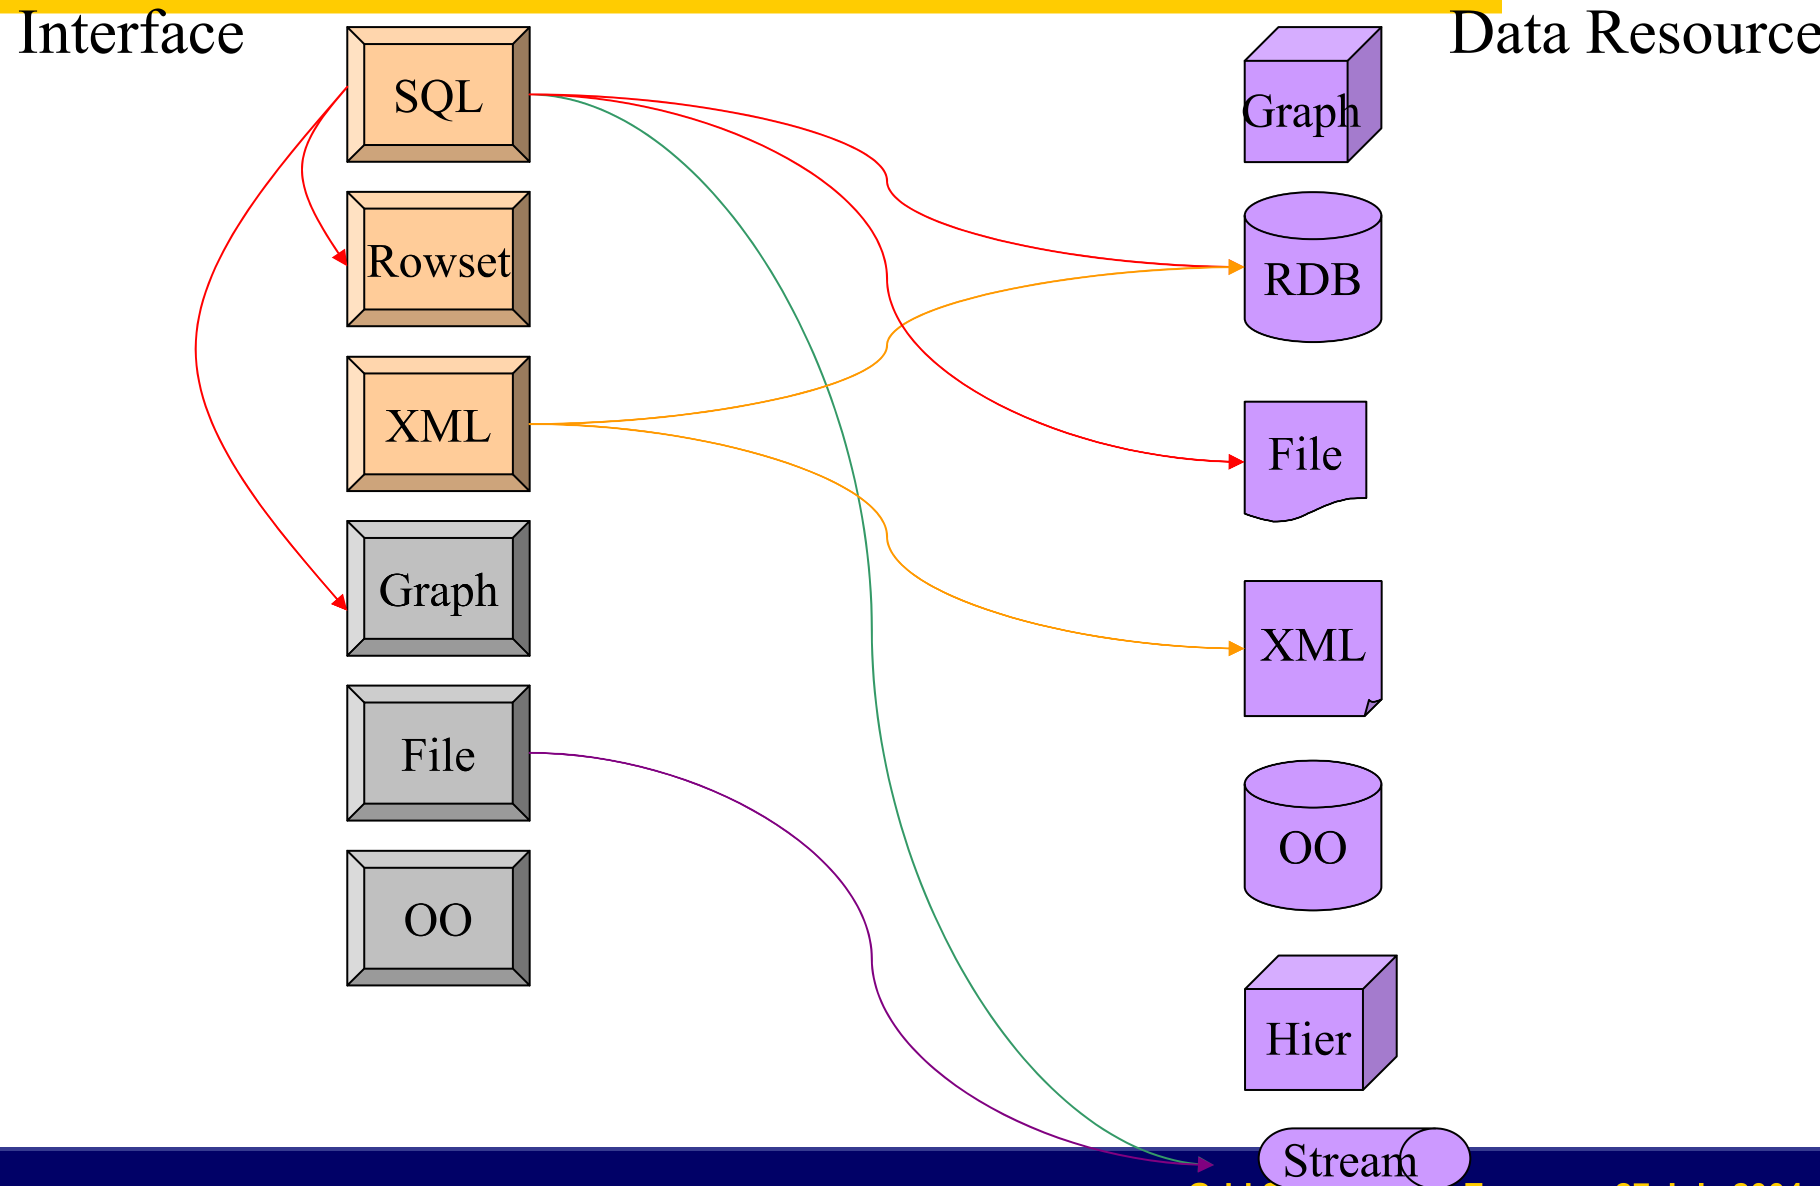

Specification Names

- **Web Services Data Access and Integration (WS-DAI)**
  - **The specification formerly known as the Grid Data Service Specification**
  - **A paradigm-neutral specification of descriptive and operational features of services for accessing data**

- **The WS-DAI Realisations**
  - **WS-DAIR: for relational databases**
  - **WS-DAIX: for XML repositories**

DAIS Specification Landscape

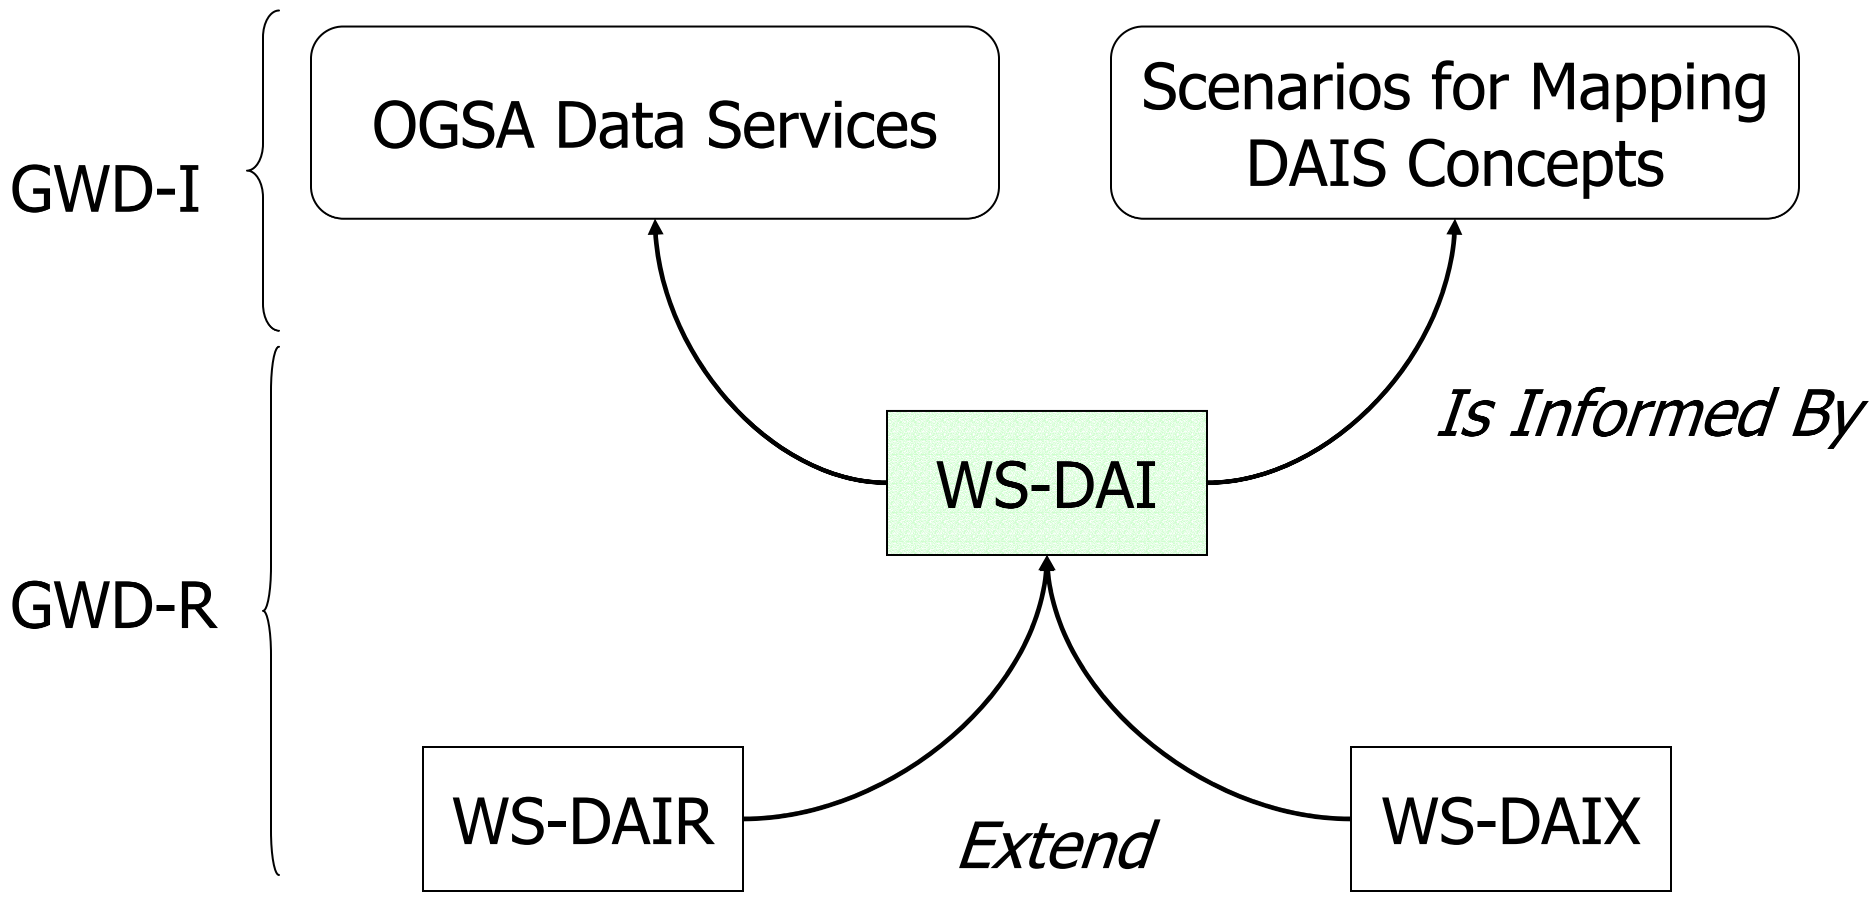

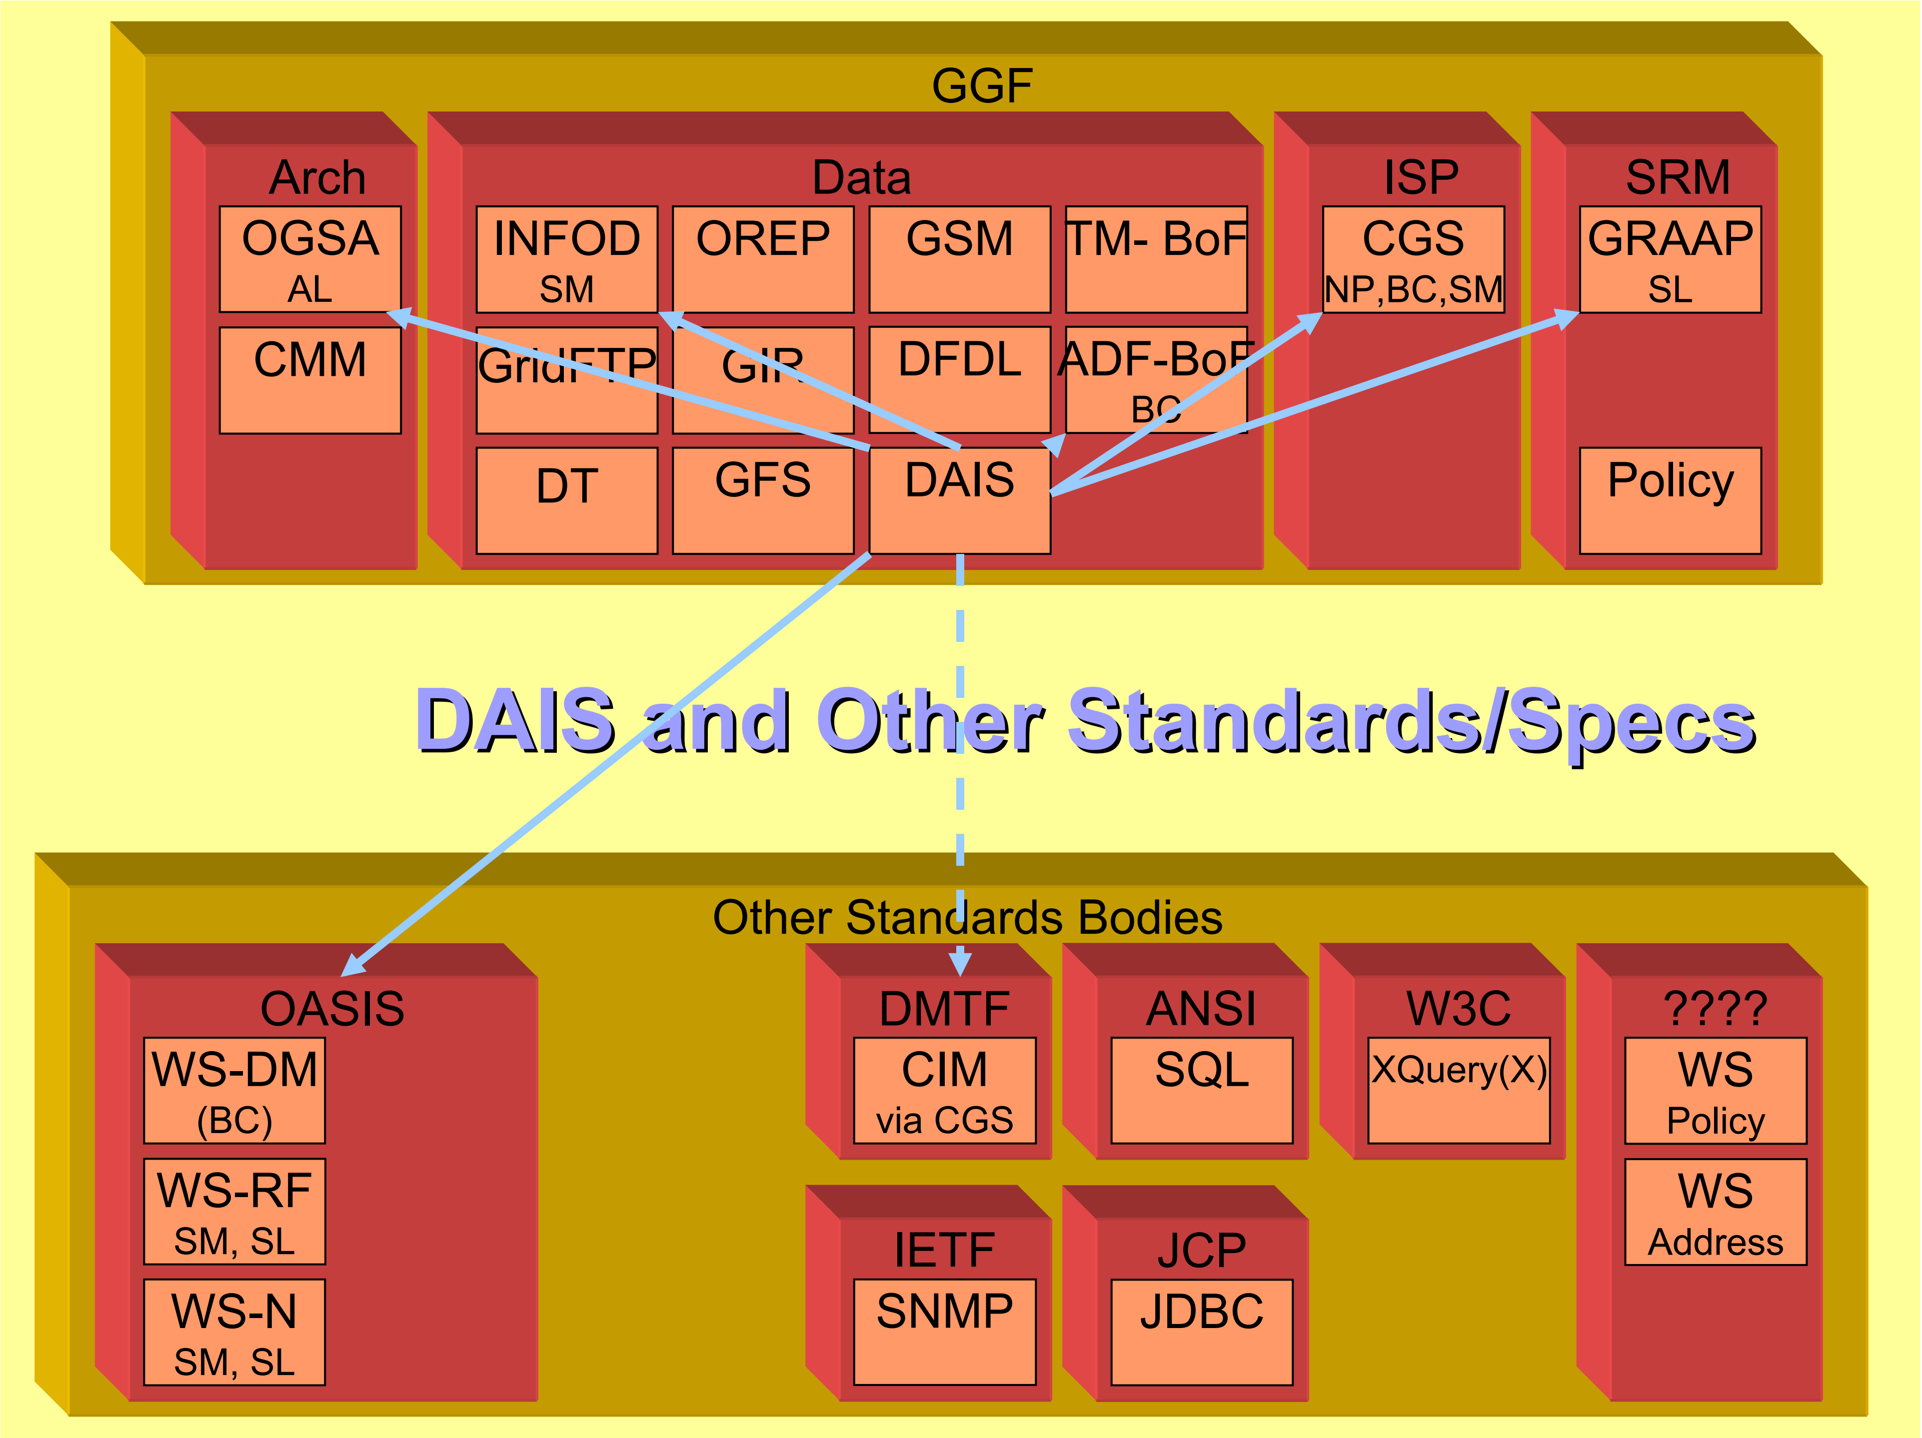
GGF
Arch
OGSA
AL
CMM
Data
INFOD
SM
OREP
GSM
TM- BoF
GridFTP
GIR
DFDL
ADF-BoF
BC
DT
GFS
DAIS
ISP
CGS
NP,BC,SM
SRM
GRAAP
SL
Policy
DAIS and Other Standards/Specs
Other Standards Bodies
OASIS
WS-DM
(BC)
WS-RF
SM, SL
WS-N
SM, SL
DMTF
CIM
via CGS
ANSI
SQL
W3C
XQuery(X)
????
WS
Policy
WS
Address
IETF
SNMP
JCP
JDBC

DAIS Data Access

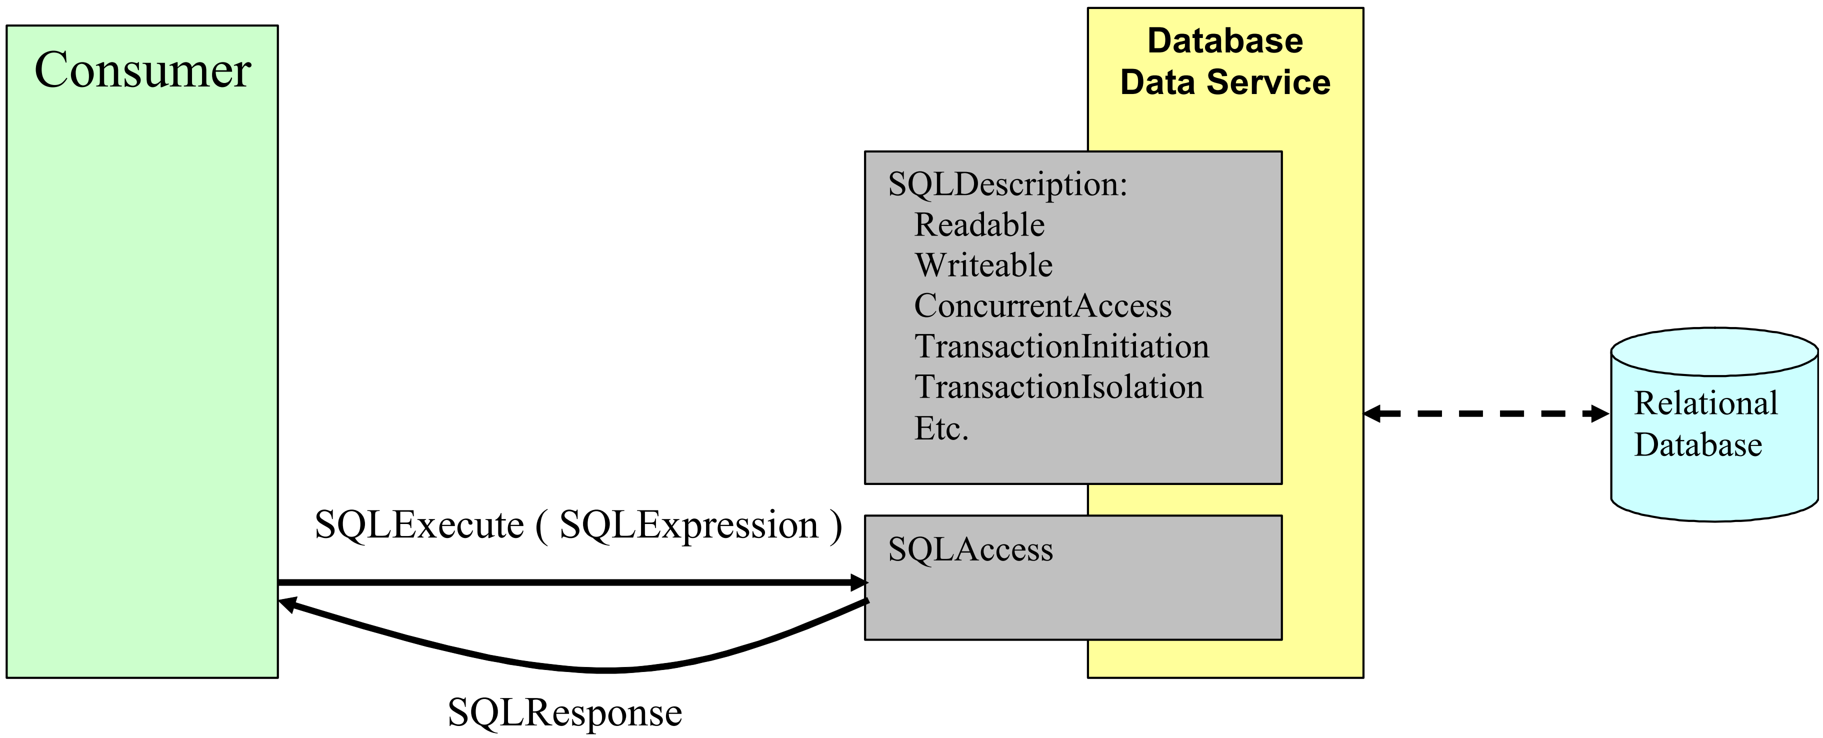

DAIS Derived Data Access

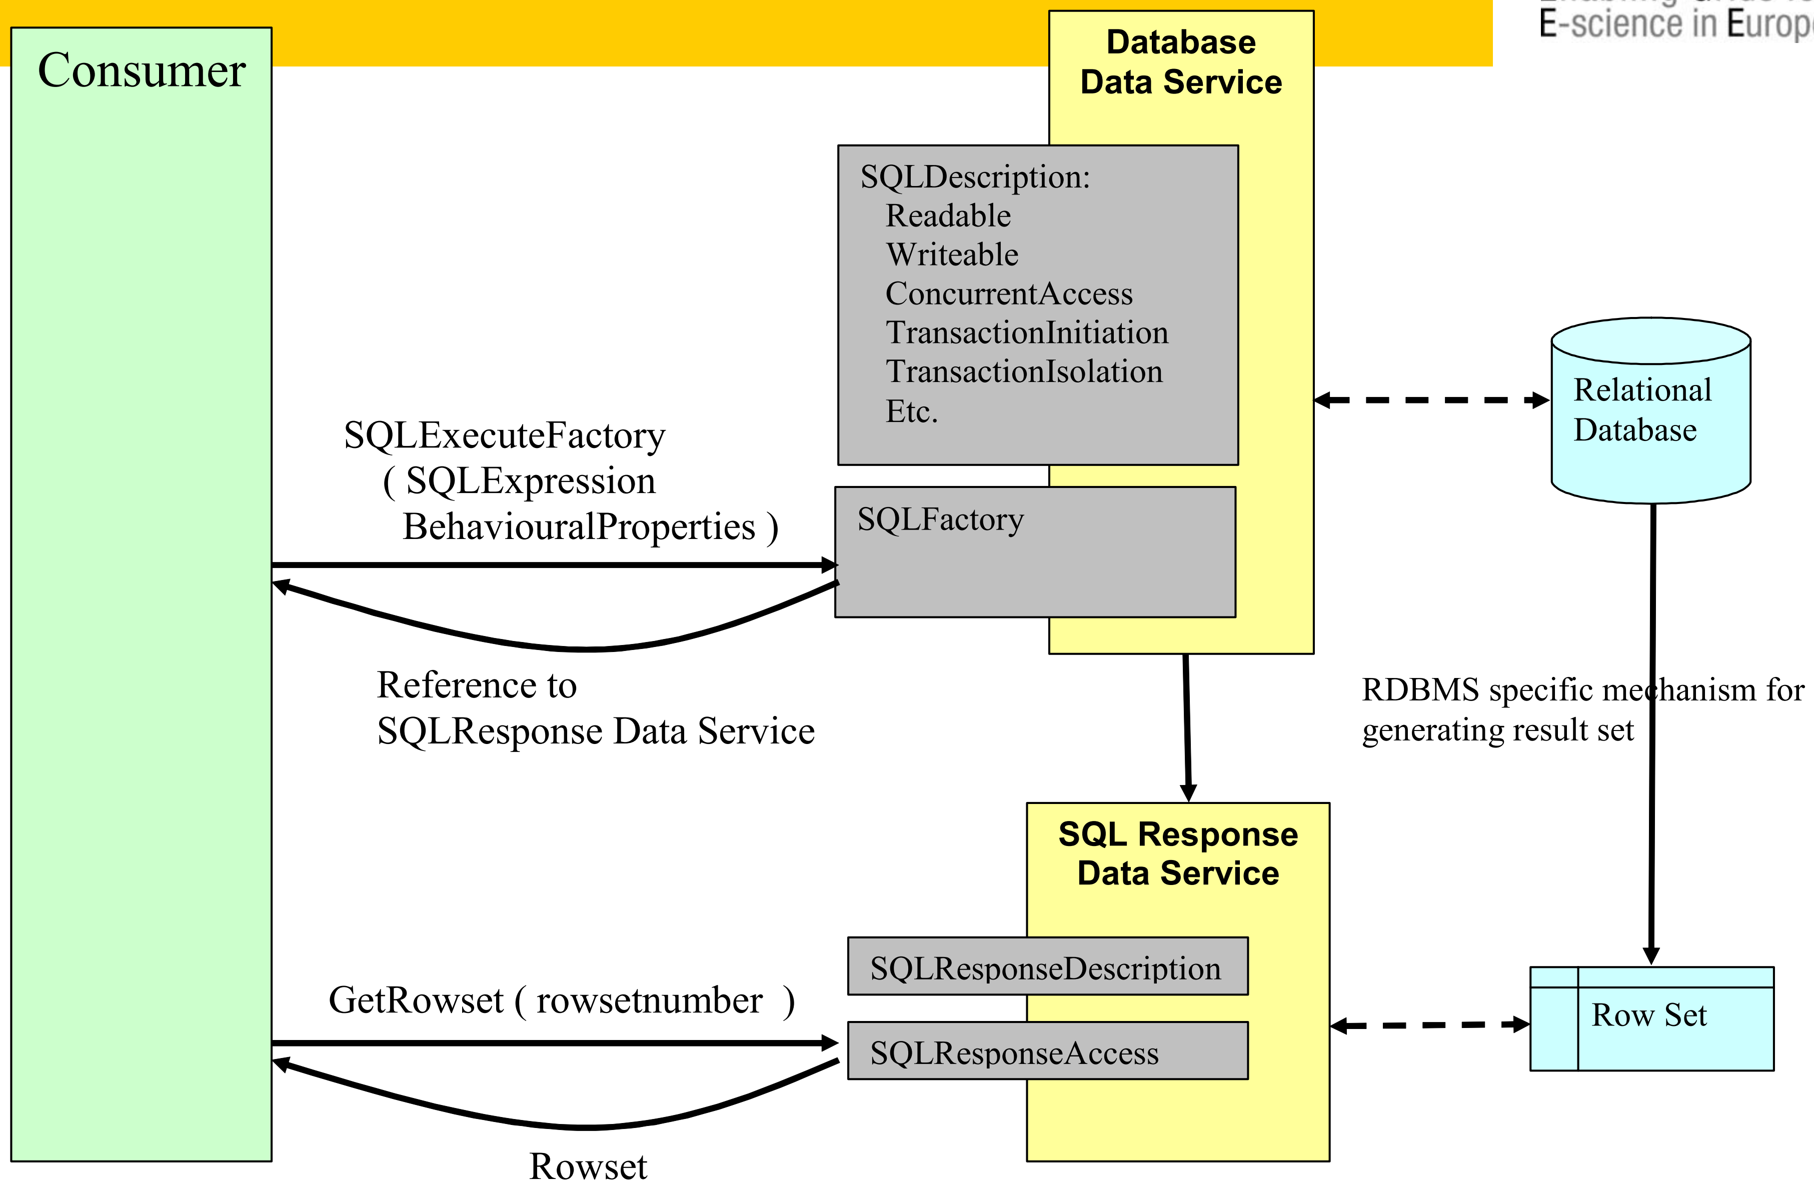

Inter-Service Data Flows

Delivery Scenarios

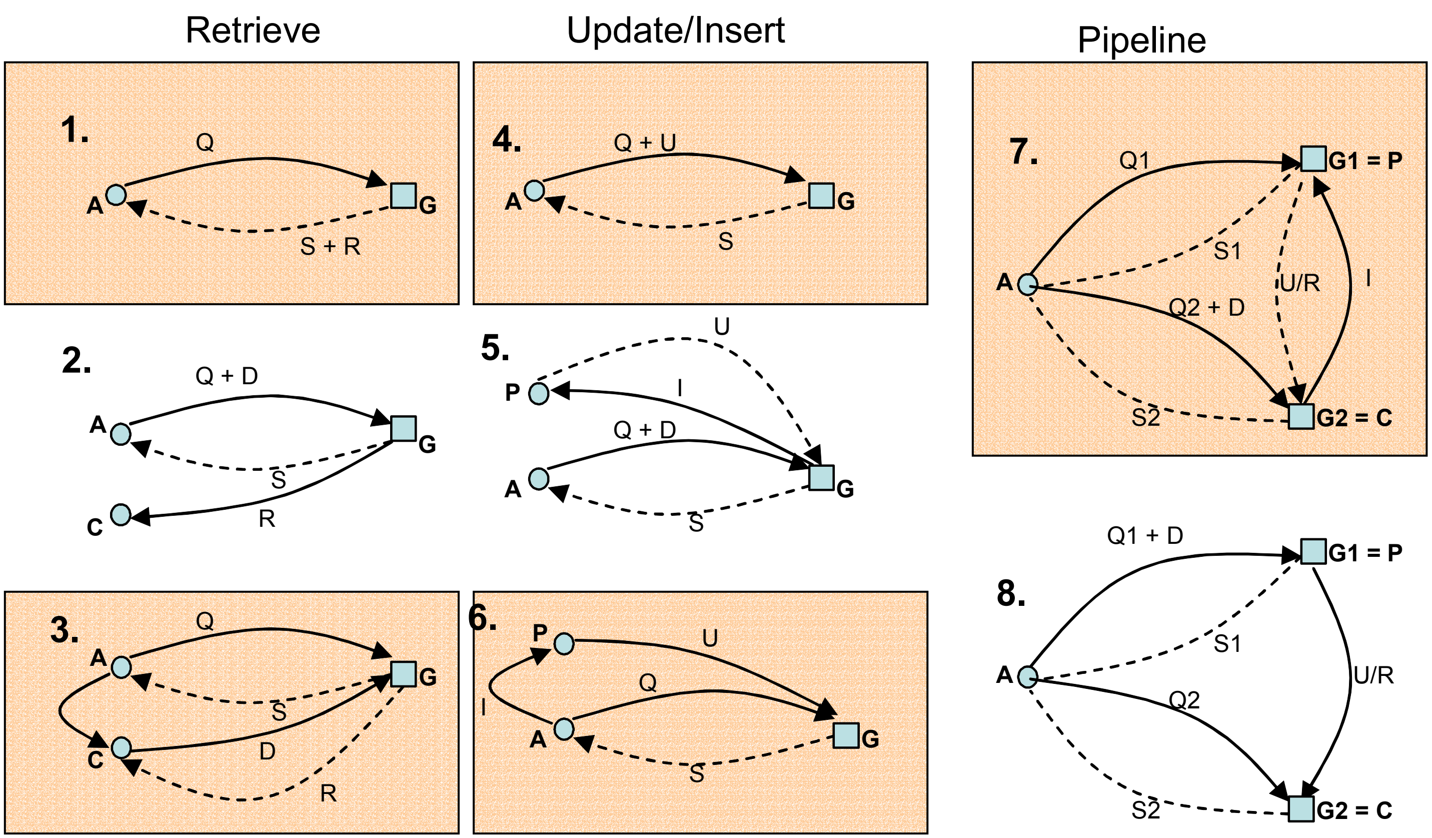

Architecture of Service Interaction

Enabling Grids for E-science in Europe

- Packaging to avoid round trips
- Unit for data movement services to handle

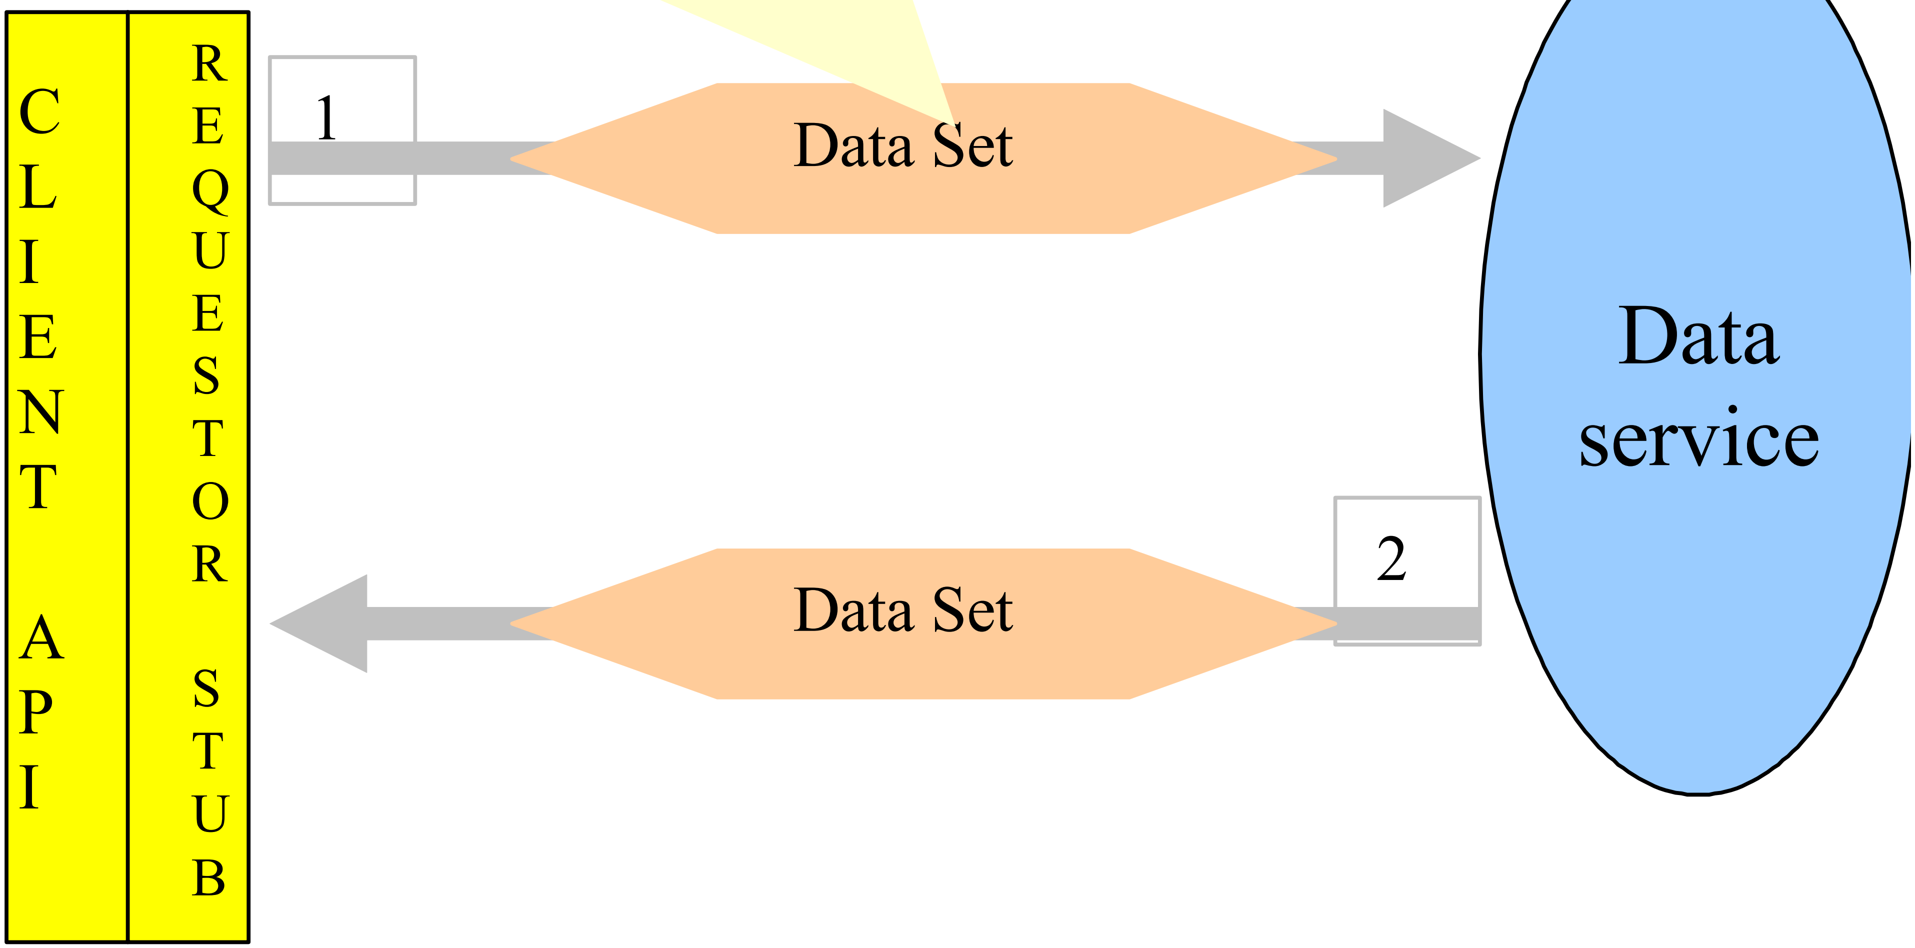

Architecture: Composing DAI

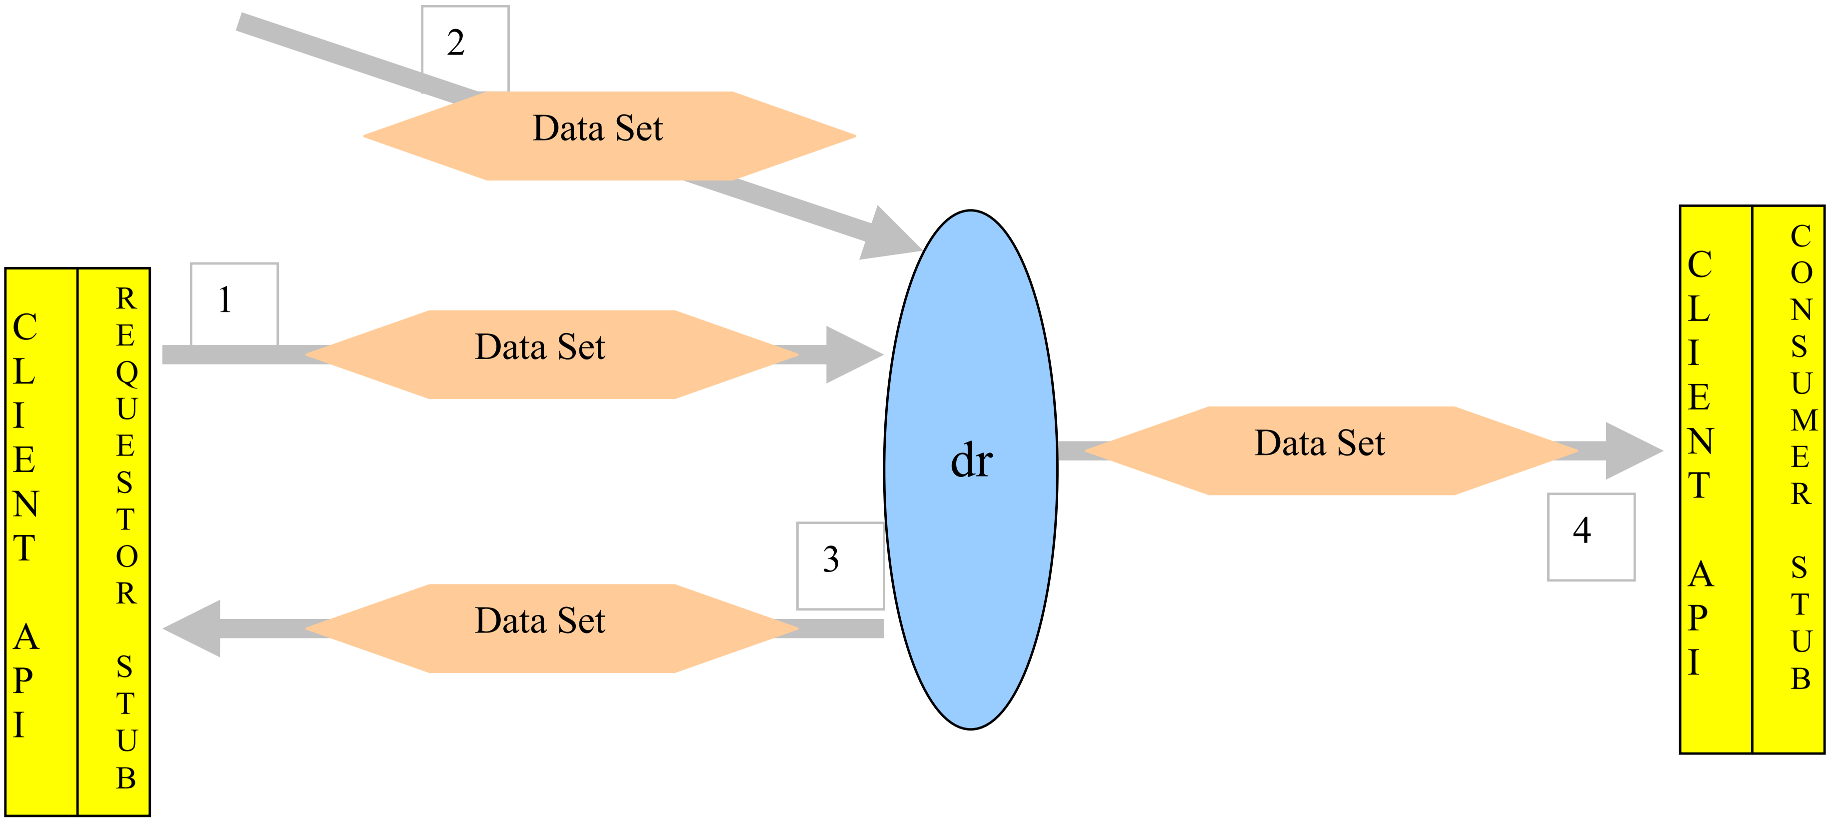

Outline

- Outline of OGSA-DAI day
- What is e-Science?
  - Collaboration & Virtual Organisations
  - Structured Data at its Foundation
- Motivation for DAI
  - Key Uses of Distributed Data Resources
  - Challenges
  - Requirements
- Standards and Architectures
  - OGSA Working Group
  - DAIS Working Group
- **Introduction to DAI**
  - **Conceptual Models**
  - **High-level Architecture**
  - **Current OGSA-DAI components**
- **Future work**

OGSA-DAI Project

OGSA-DAI

First steps towards a generic framework for integrating data access and computation

Using the grid to take specific classes of *computation nearer to the data*

Kit of parts for building tailored access and integration applications

Investigations to inform DAIS-WG

One reference implementation for DAIS

Releases publicly available NOW

OGSA-DAI Partners

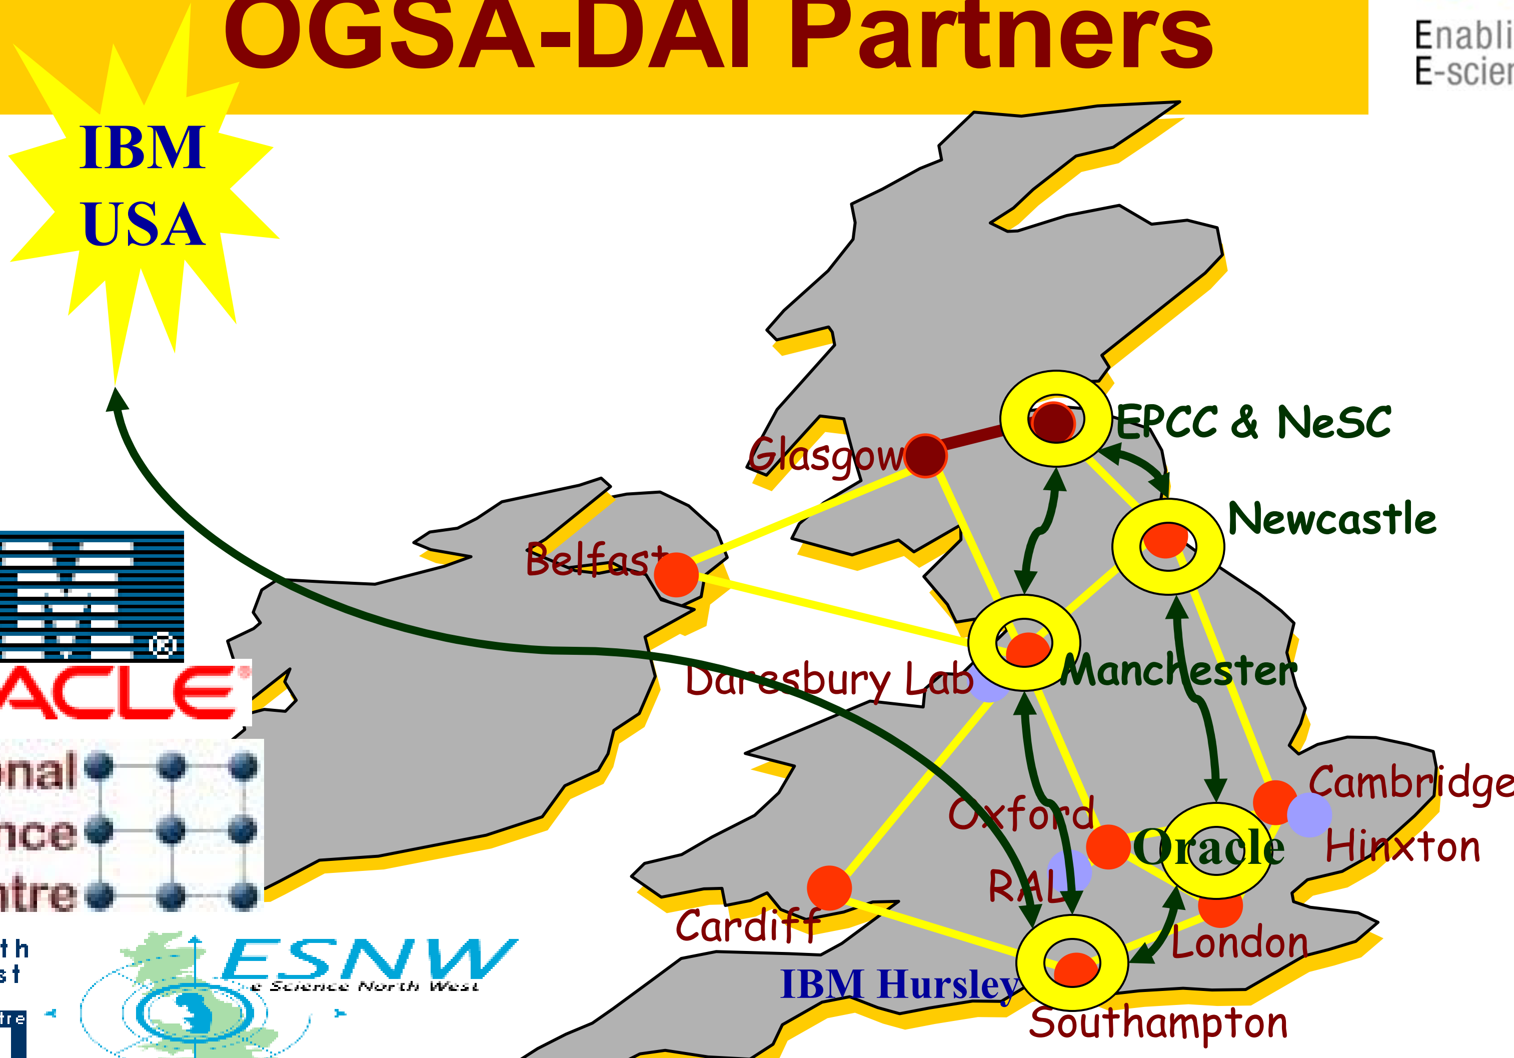

$5 million, 20 months, started February 2002

Additional 24 months, started October 2003

Infrastructure Architecture

Enabling Grids for E-science in Europe

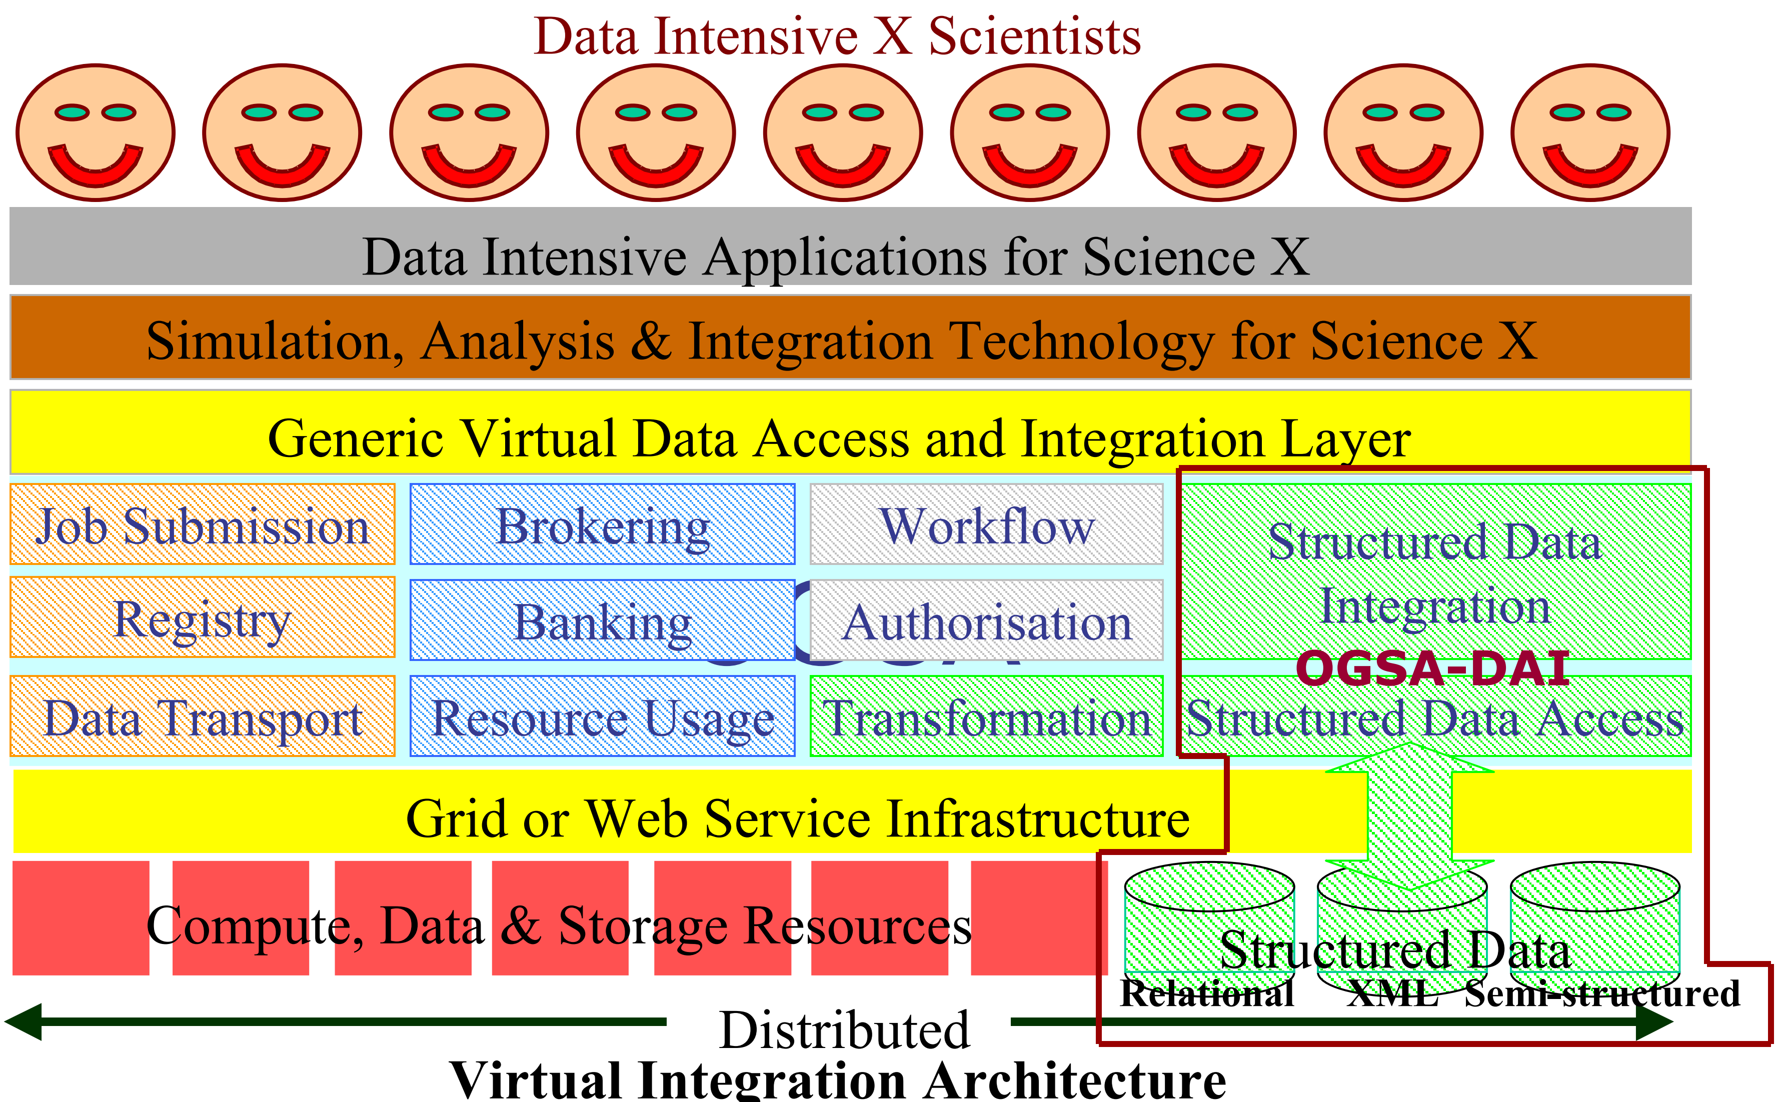

OGSA-DAI Product

Web Service Architecture

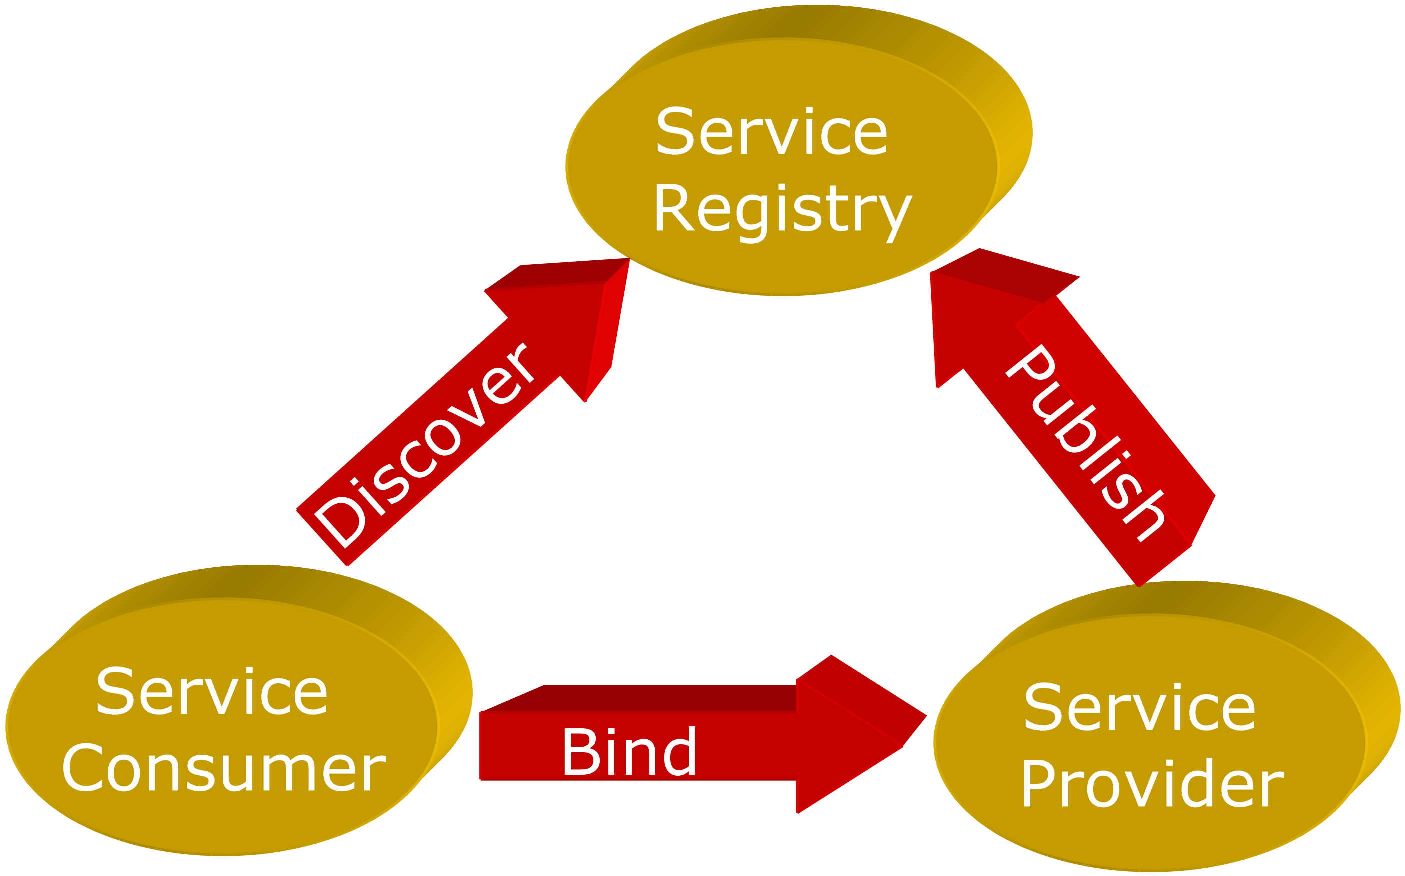

OGSA-DAI Service Architecture

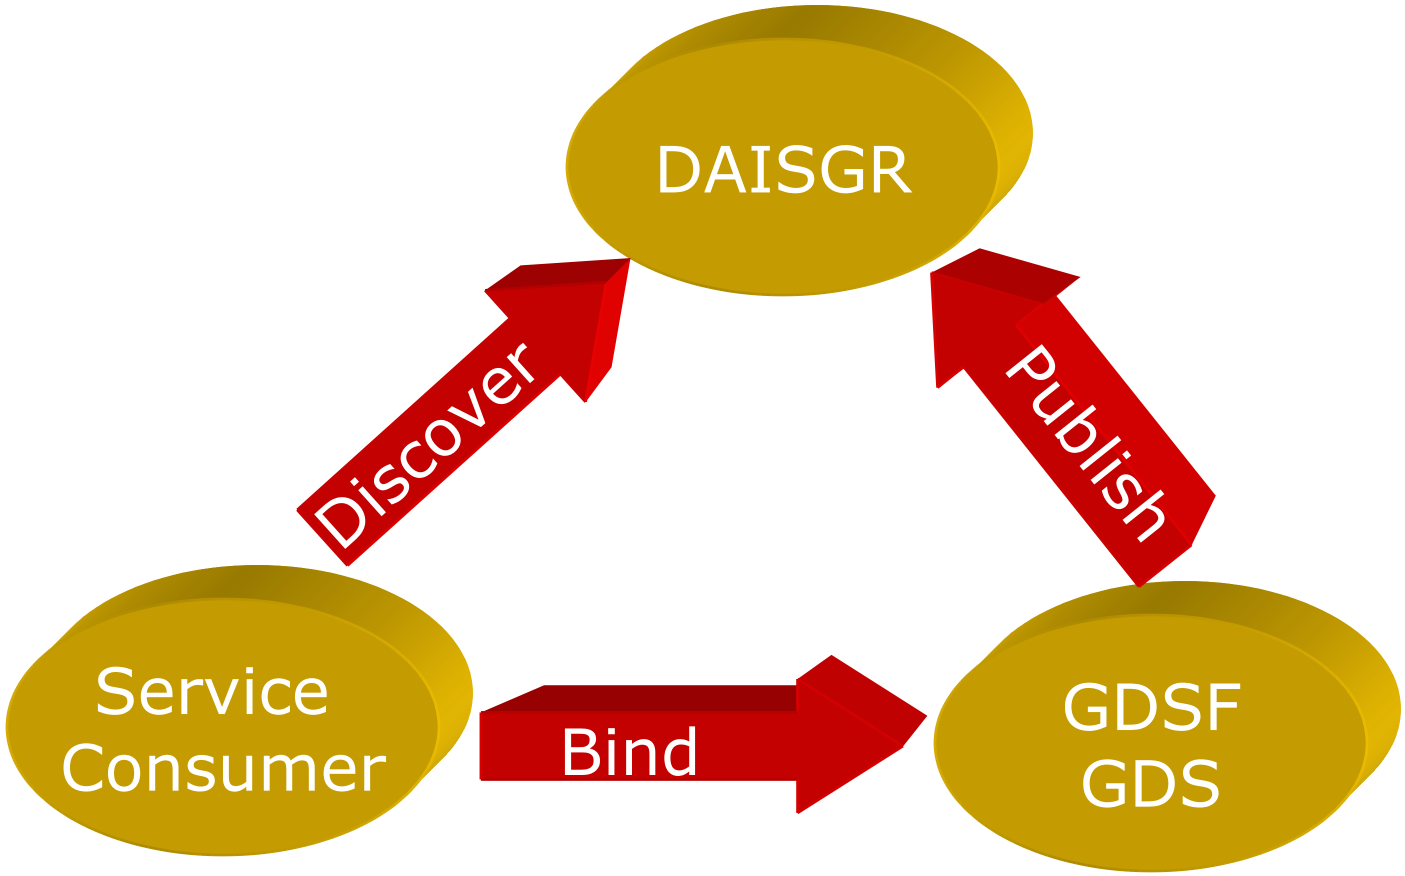

OGSA-DAI Services

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

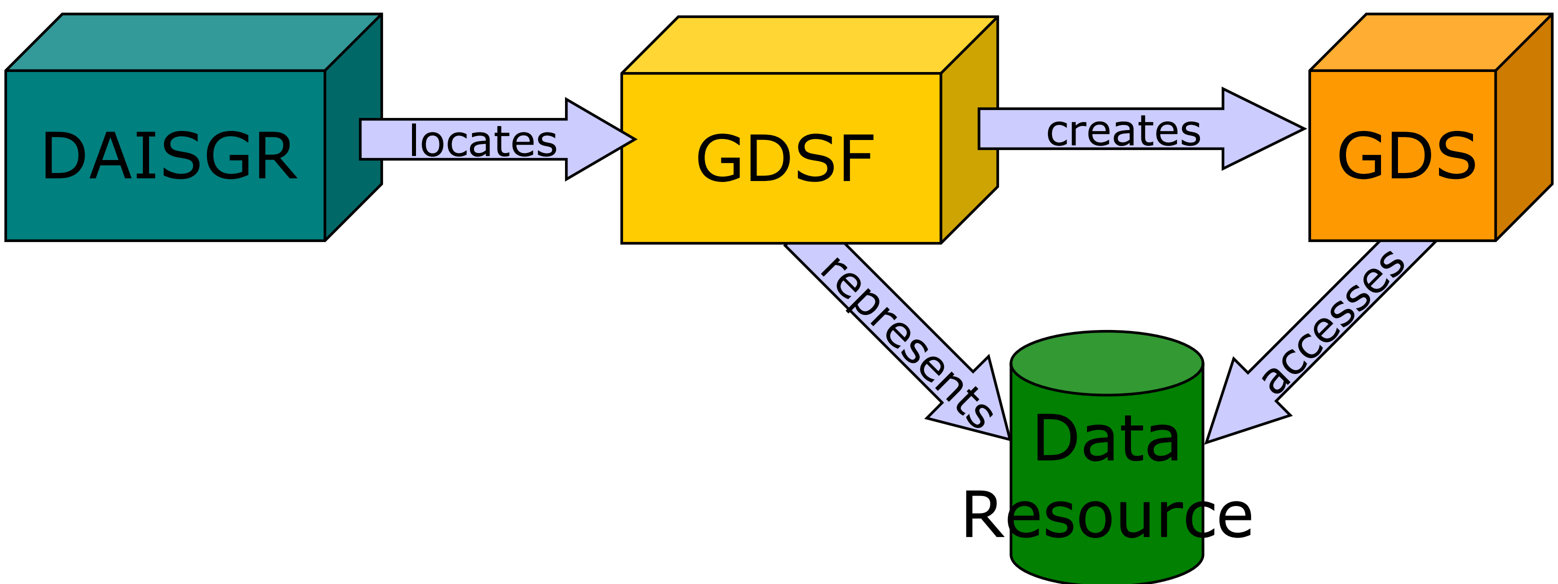

GDSF and GDS

- **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**

DAISGR

- **DAI Service Group Registry (DAISGR)**
  - **Persistent service**
  - **Based on OGSI ServiceGroups**
  - **GDSFs may register with DAISGR**
  - **Clients access DAISGR to discover**
    - **Resources**
    - **Services (may need specific capabilities)**
      - **Support a given portType or activity**

Interaction Model: Start up

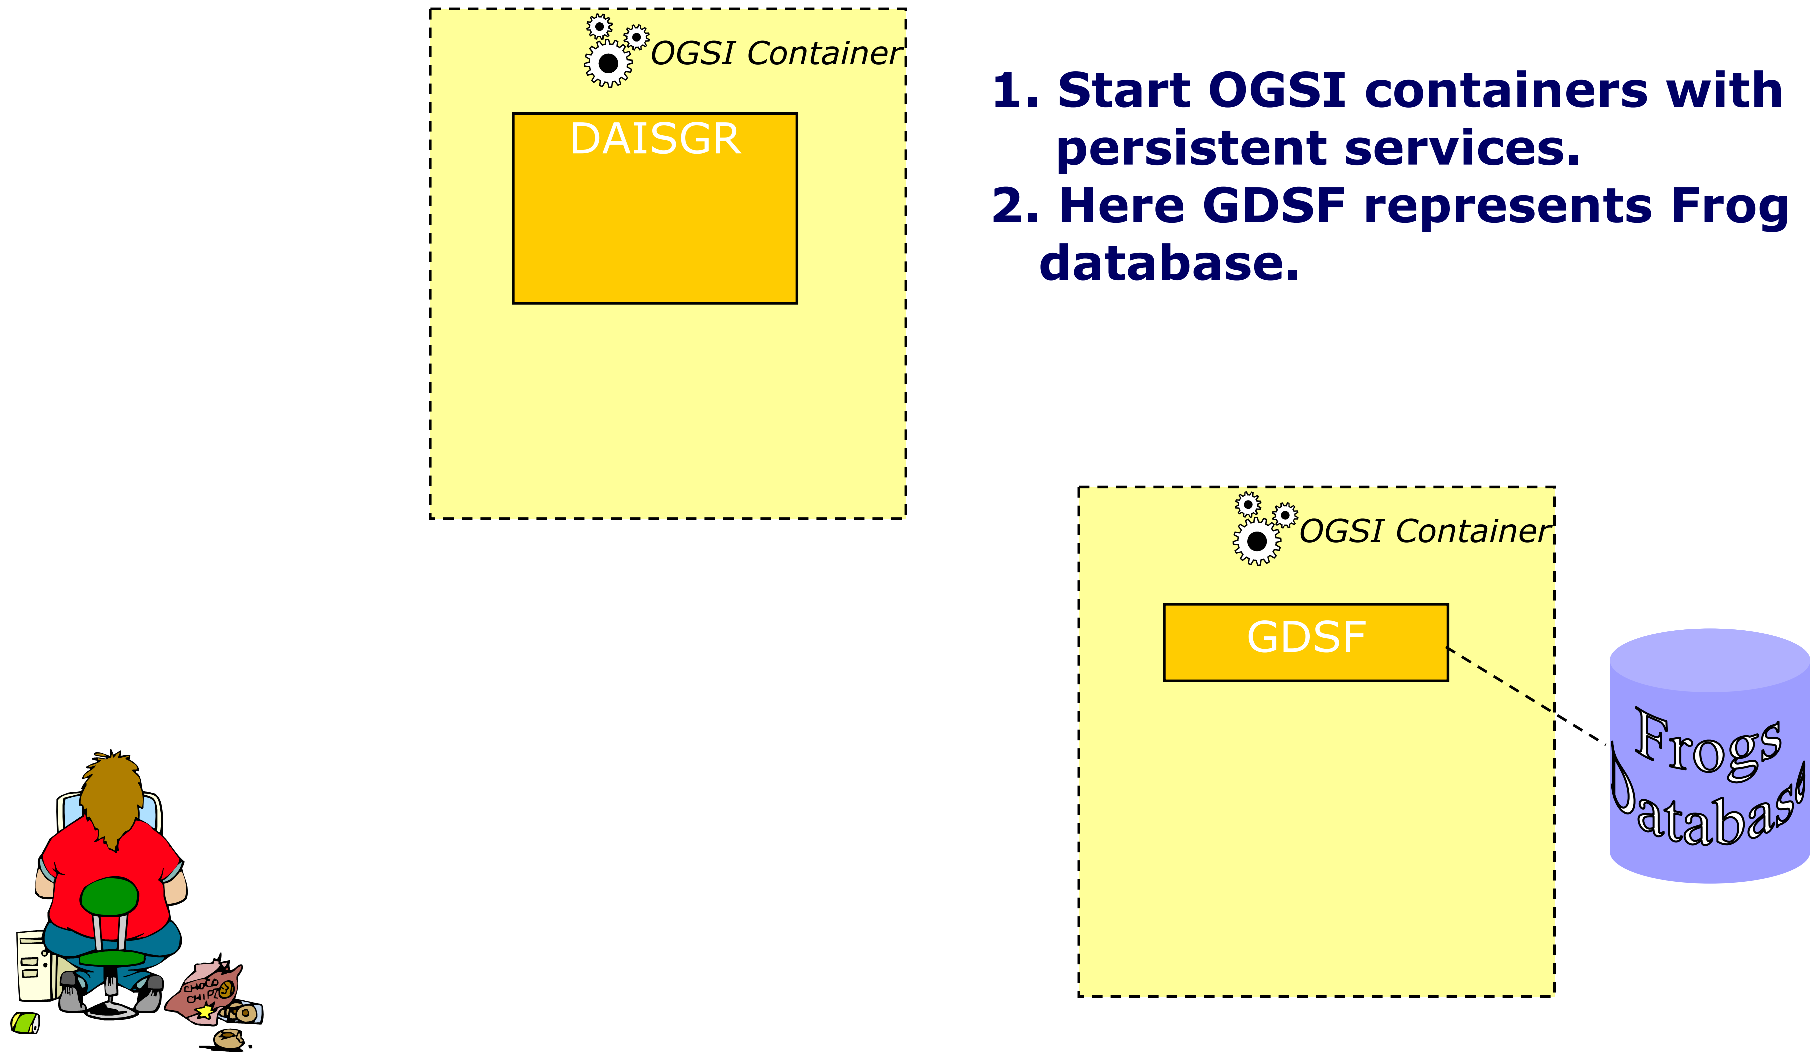

1. Start OGSI containers with persistent services.
2. Here GDSF represents Frog database.

Interaction Model: Registration

Enabling Grids for E-science in Europe

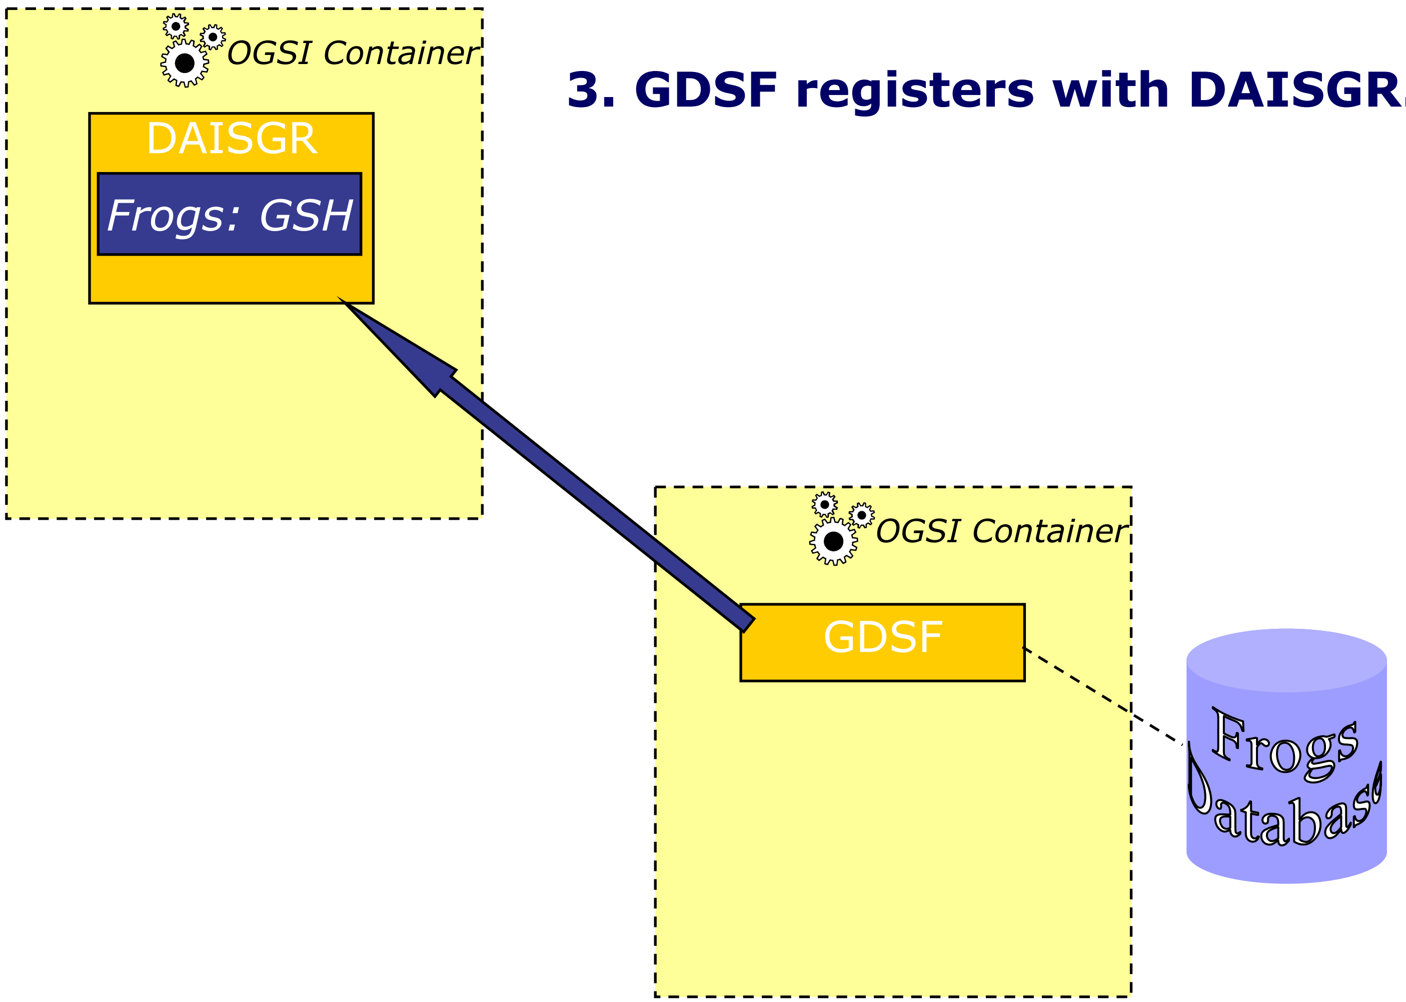

3. GDSF registers with DAISGR.

Interaction Model: Discovery

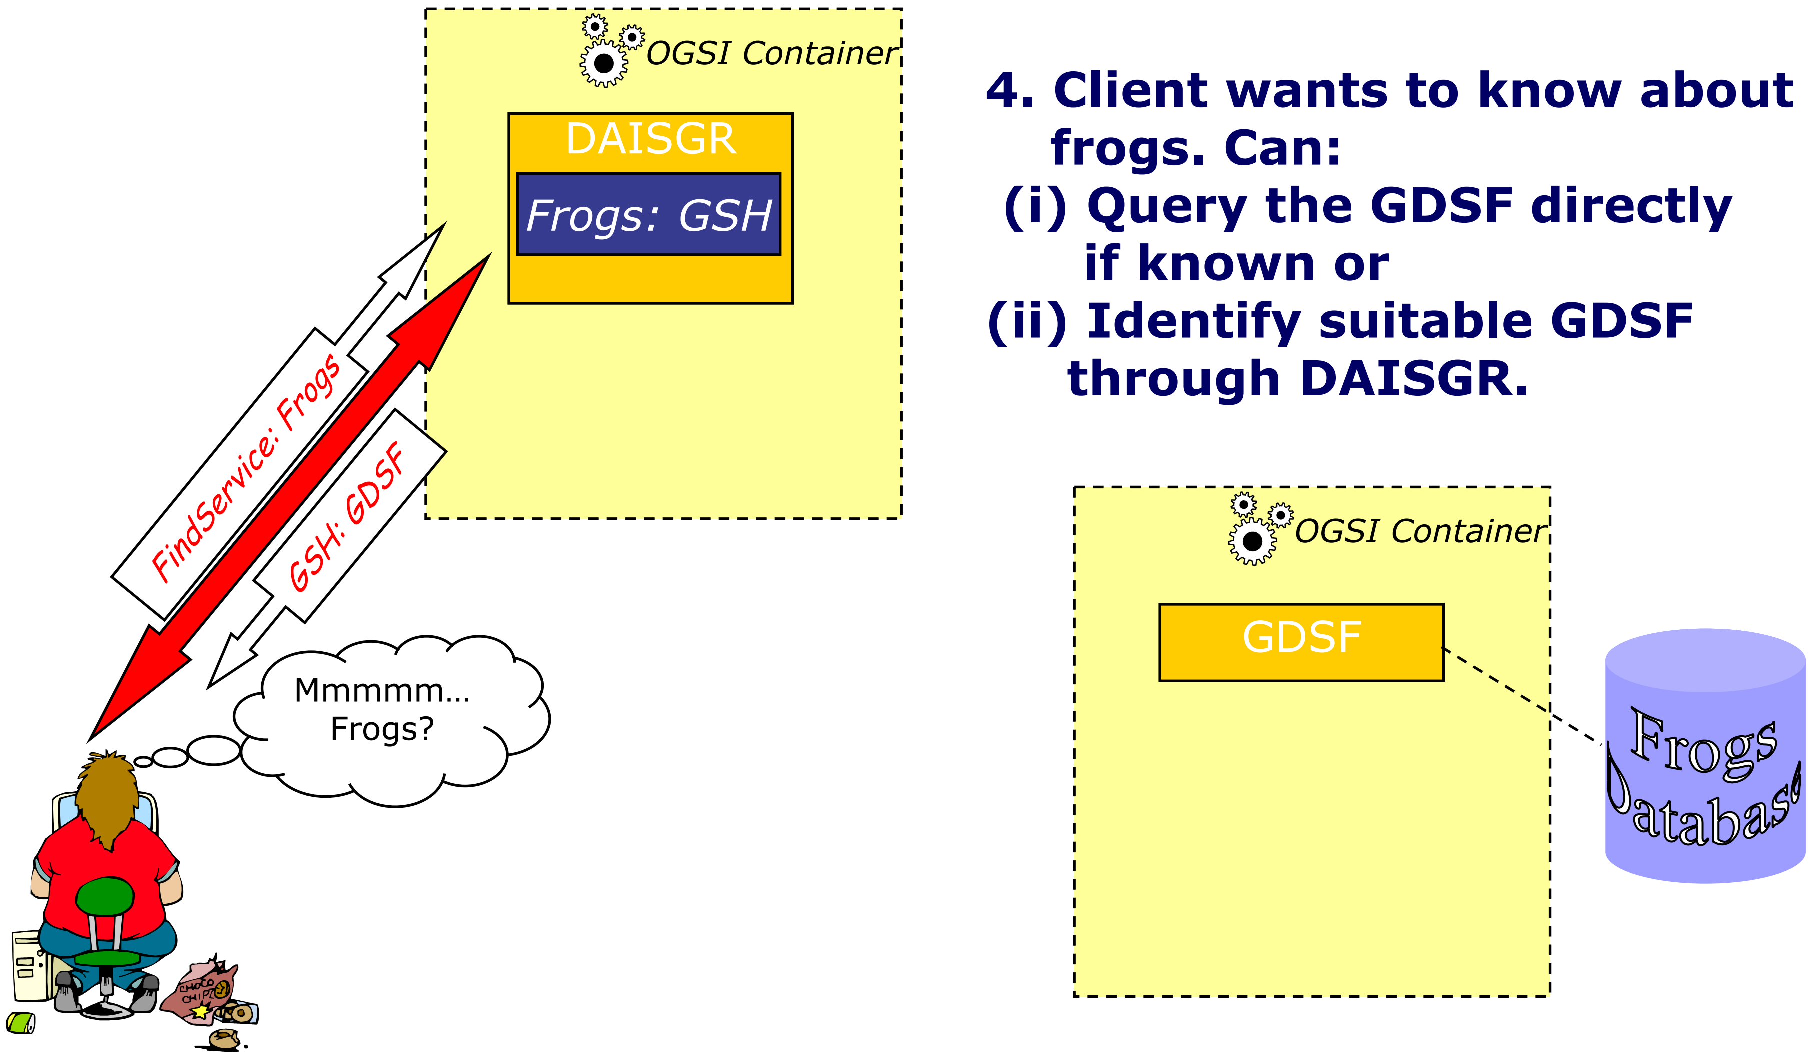

4. Client wants to know about frogs. Can:
(i) Query the GDSF directly if known or
(ii) Identify suitable GDSF through DAISGR.

Interaction Model: Service Creation

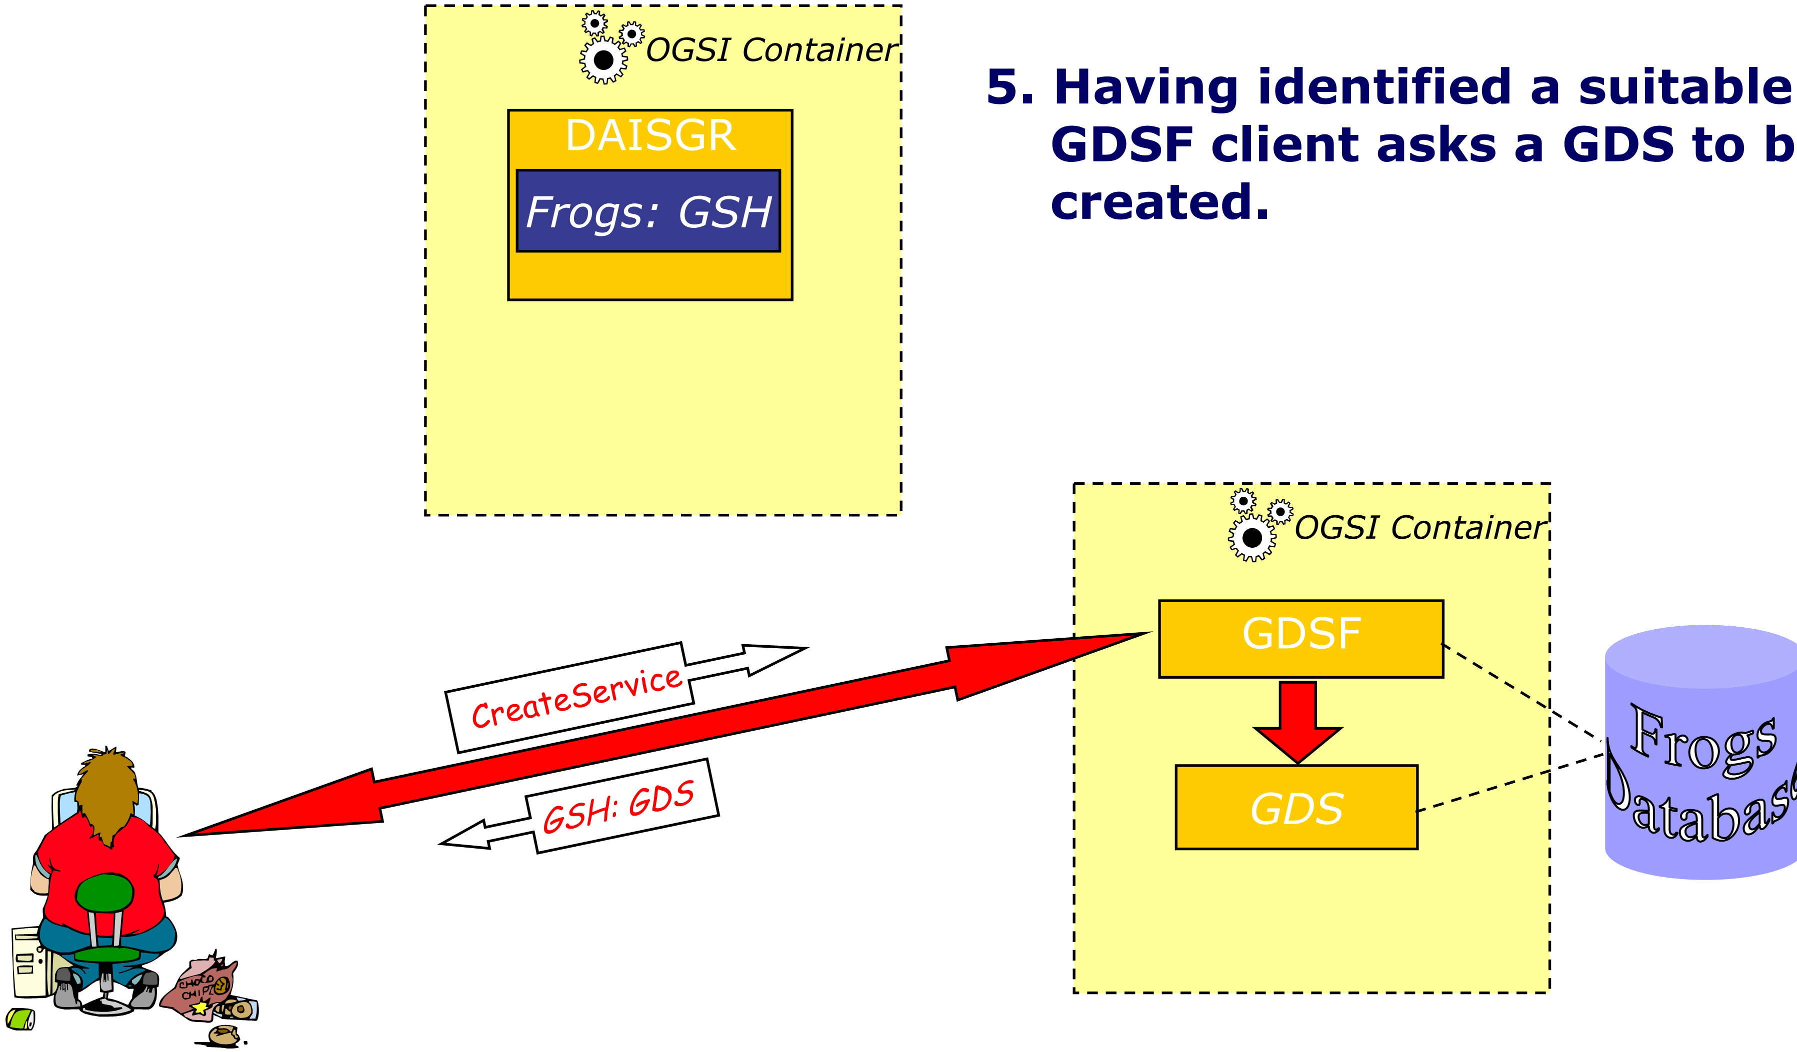

5. Having identified a suitable GDSF client asks a GDS to be created.

Interaction Model: Perform

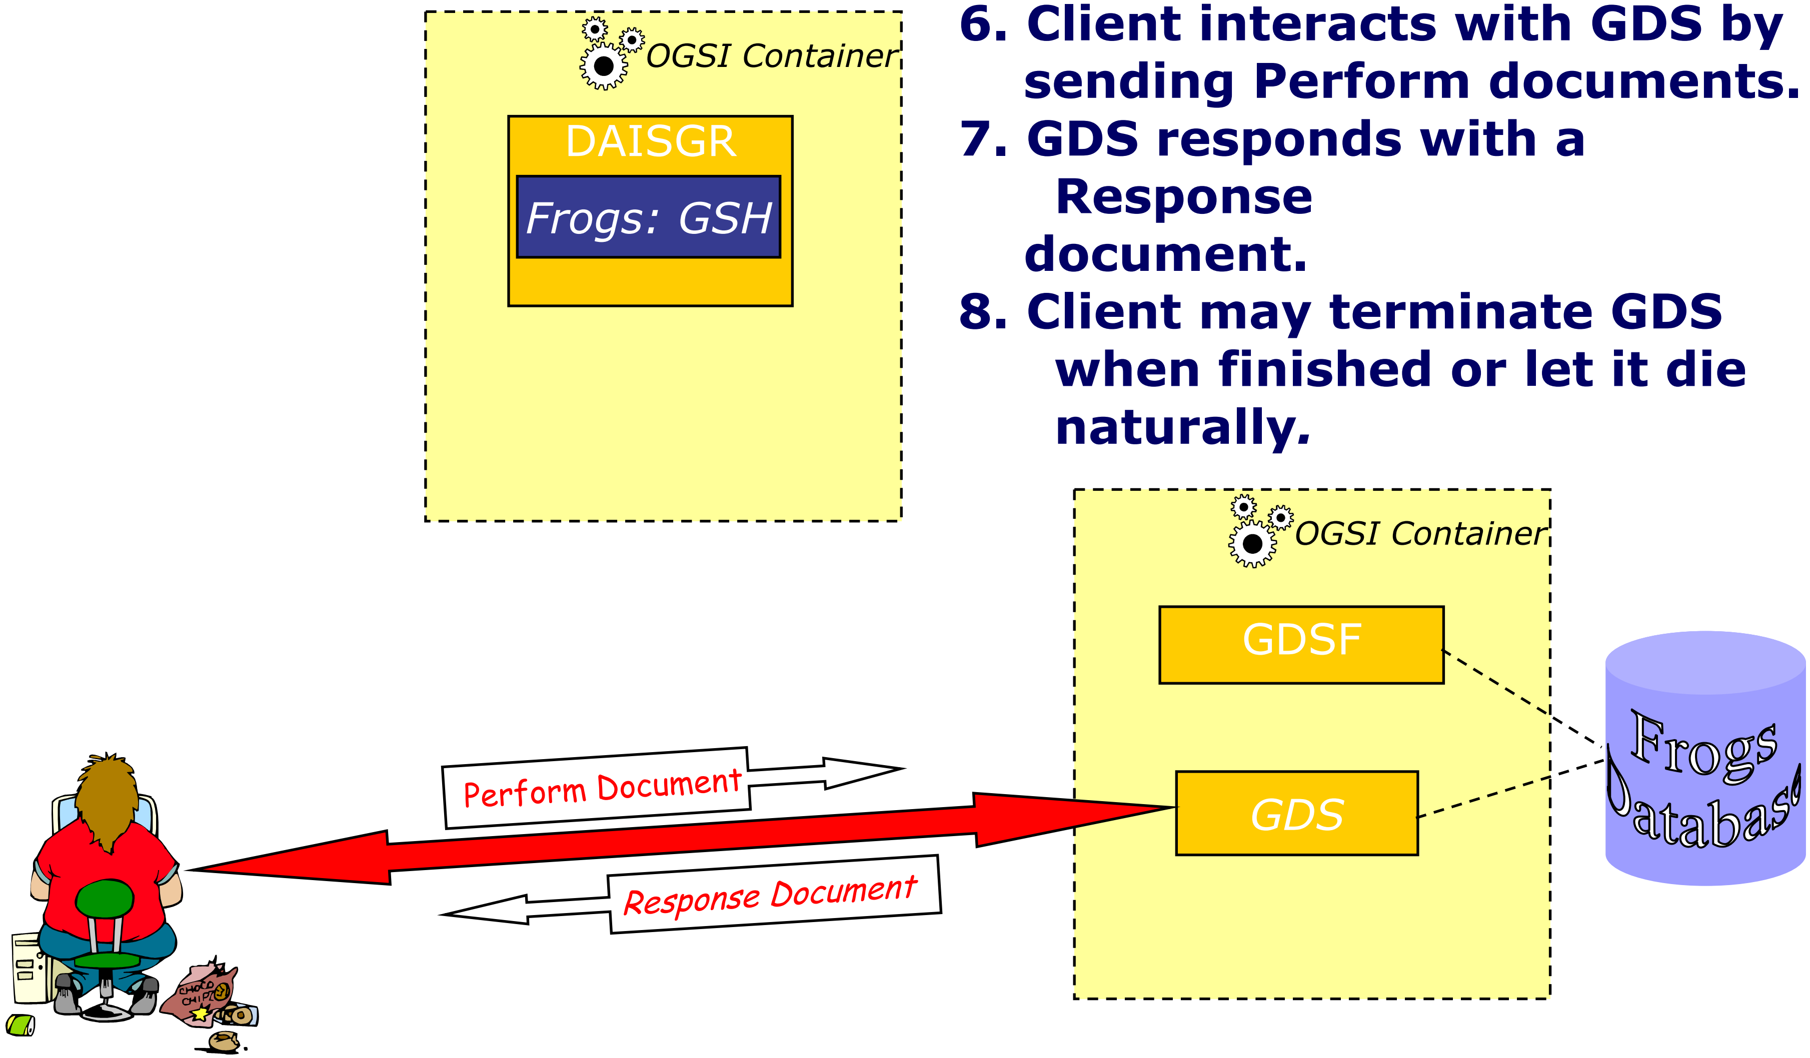

6. Client interacts with GDS by sending Perform documents.
7. GDS responds with a Response document.
8. Client may terminate GDS when finished or let it die naturally.

Interaction Model: Sum up

- **Only described 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 that OGSA-DAI is doing ...**

Data Access & Integration Services

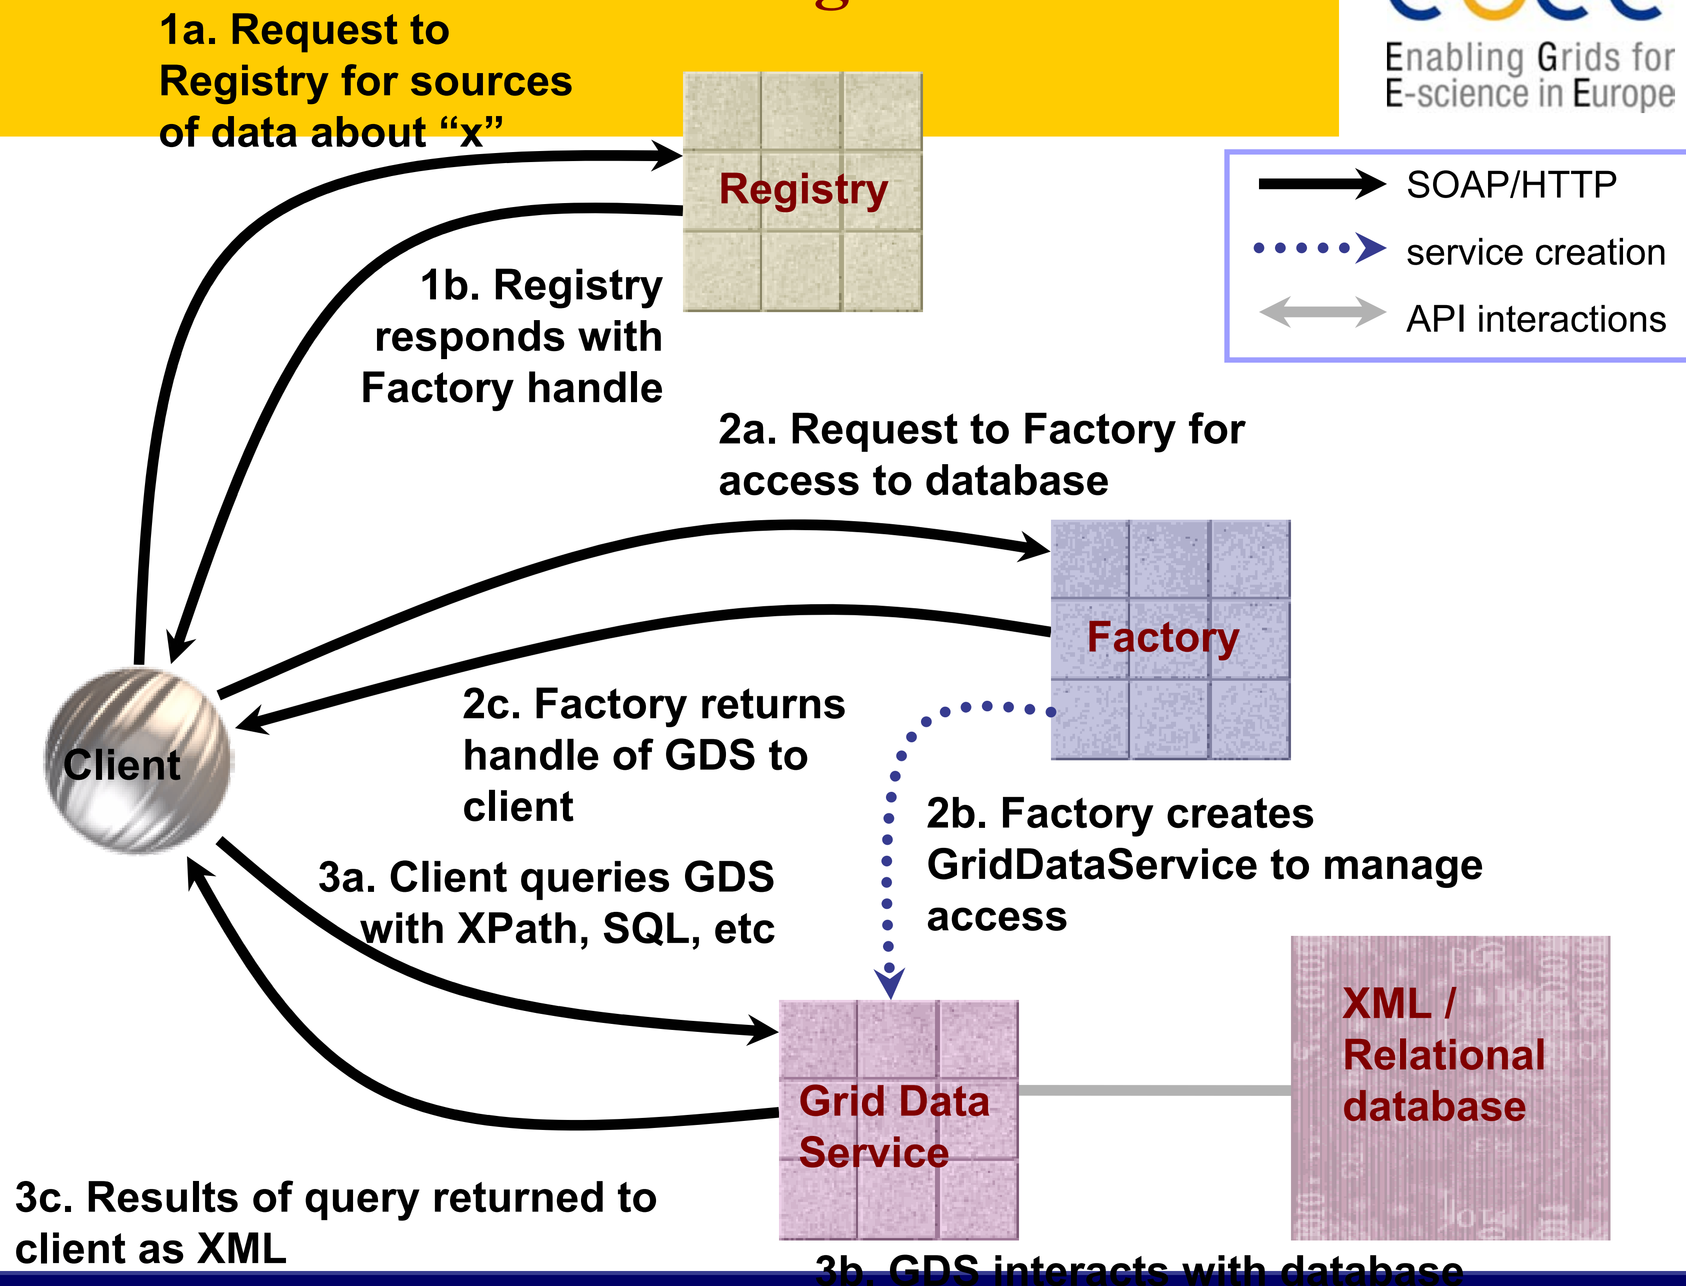

Future Work

Future DAI Services – Integrated in Tools

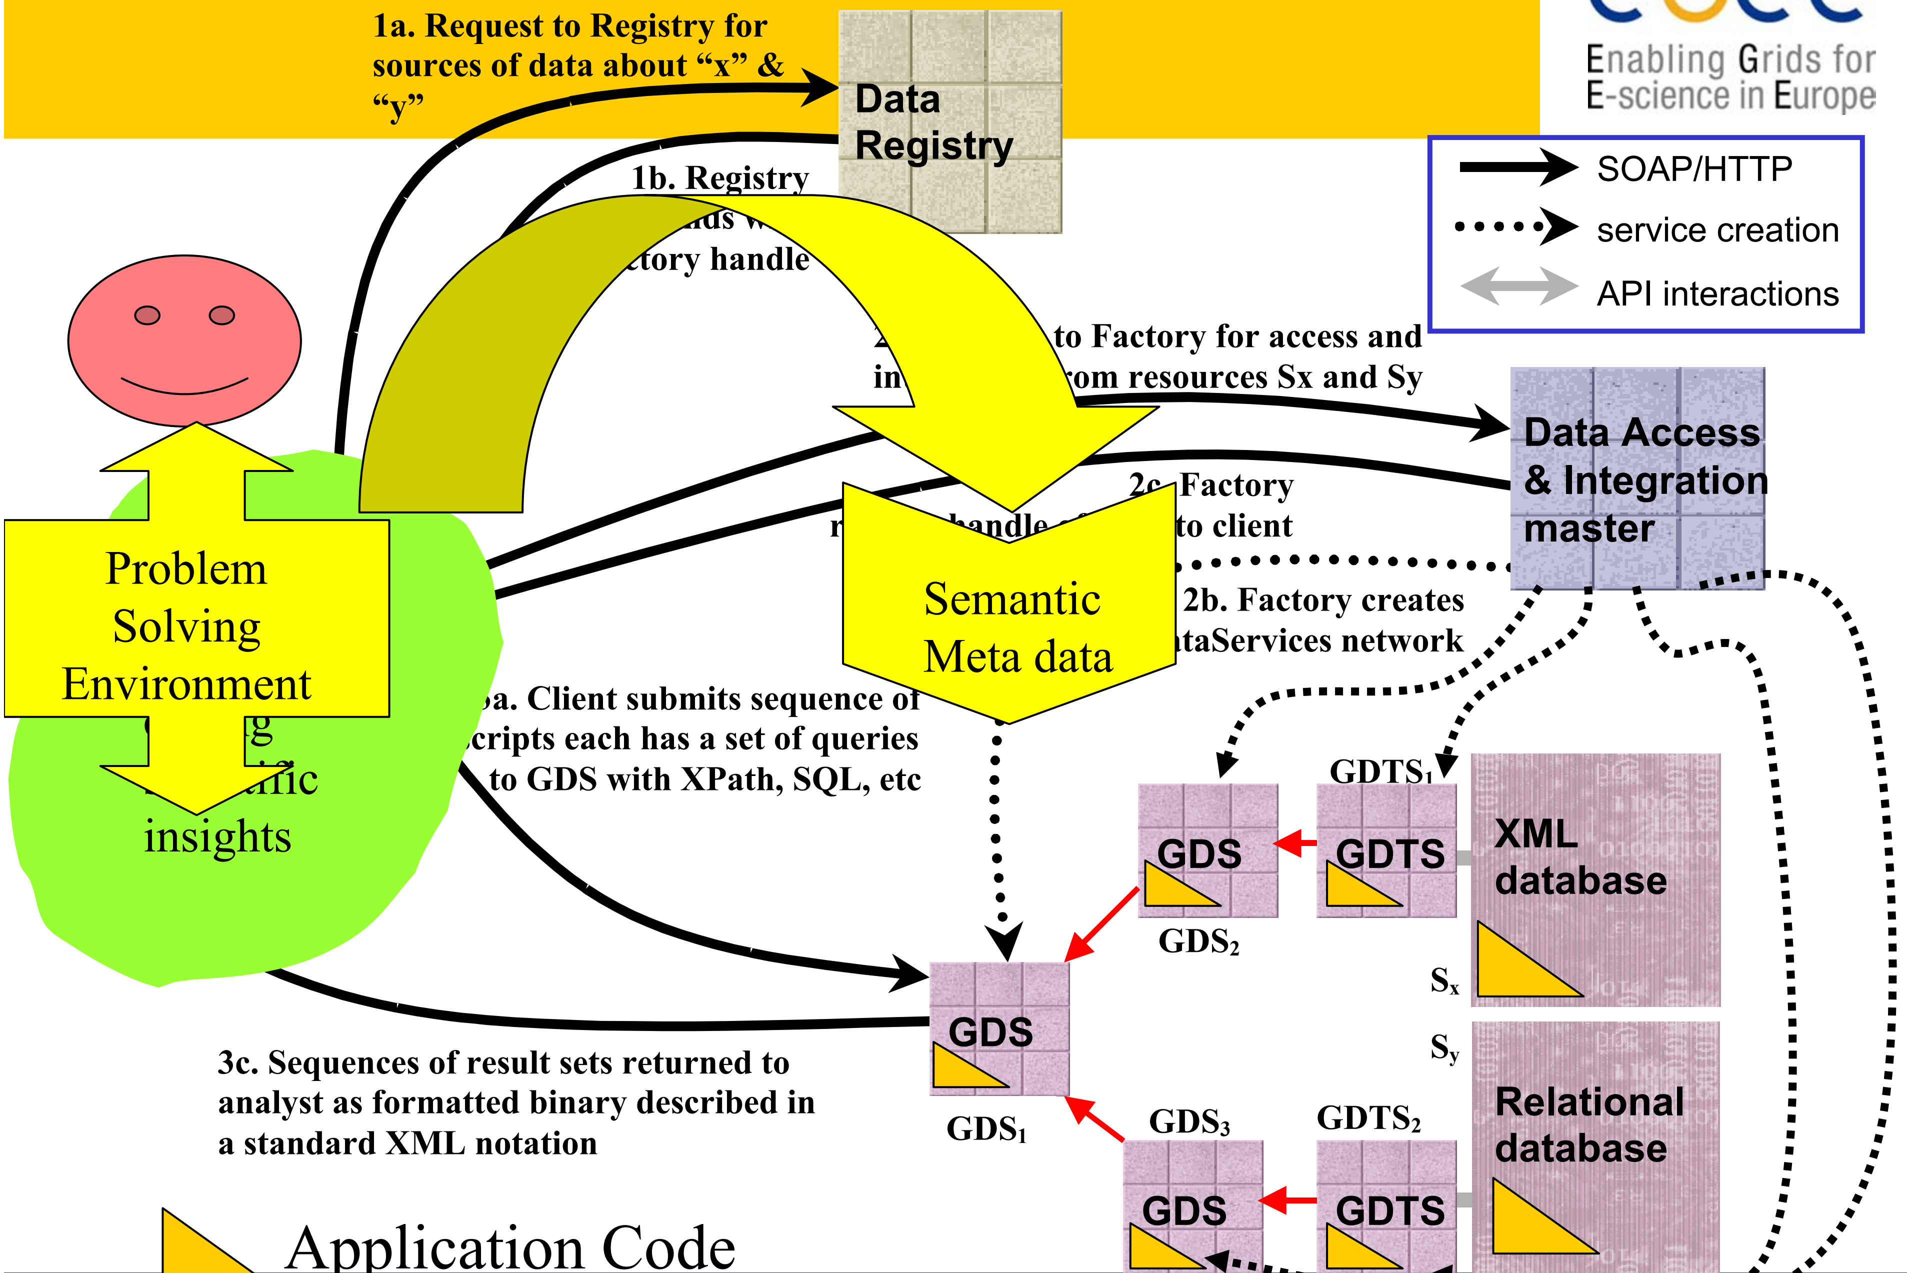

Extensibility a Necessity

Enabling Grids for E-science in Europe

- **Data resources**
  - Unbounded variety
- **Data access languages**
  - Established standards
    - With many variants
  - SQL, OQL, semi-structured query, domain languages
- **Investment in DBs, DBMSs, File Stores, Bulk stores, …**
  - Not sensible to expect them to change to fit us
- **Data Access Models must be extensible**
  - Static extension used extensively by OGSA-DAI users

Should extensibility be supported by foundation interfaces?

Move Computation to Data

- **Code scale**
  - **Depends on wet-ware**
    - **No noticeable rate of improvement**
- **Data scale**
  - **Grows Moore's Law or Moore's Law²**
- **Analysis of data**
  - **Extracts & derivatives used**
    - **Often smaller – more value for current inv**
- **Implies move code to data**
  - **SQL, Xquery, Java code, DB Procs, Dynamic DB procs, …**
- **Extensibility mechanisms used by OGSA-DAIers**
- **Java mobility (e.g. DataCutter), database procedures, …**

Increasingly necessary

Application control or higher-level service decisions

Integration is Everything

- No business or research team is satisfied with one data resource
- Domain-specialist driven
  - Dynamic specification of combinatio
  - Iterative processes – range of time s
- Sources inevitably heterogeneou
  - Content, structure & policies time-va
- Robust & stable steerable integr
  - Higher-level services over multiple r
  - Fundamental requirements for (re)n
- Integrate Data Handling with Co

Federation or Virtualisation preceding integration or kit of integration tools to be interwoven with an application?

Multiple tasks / request

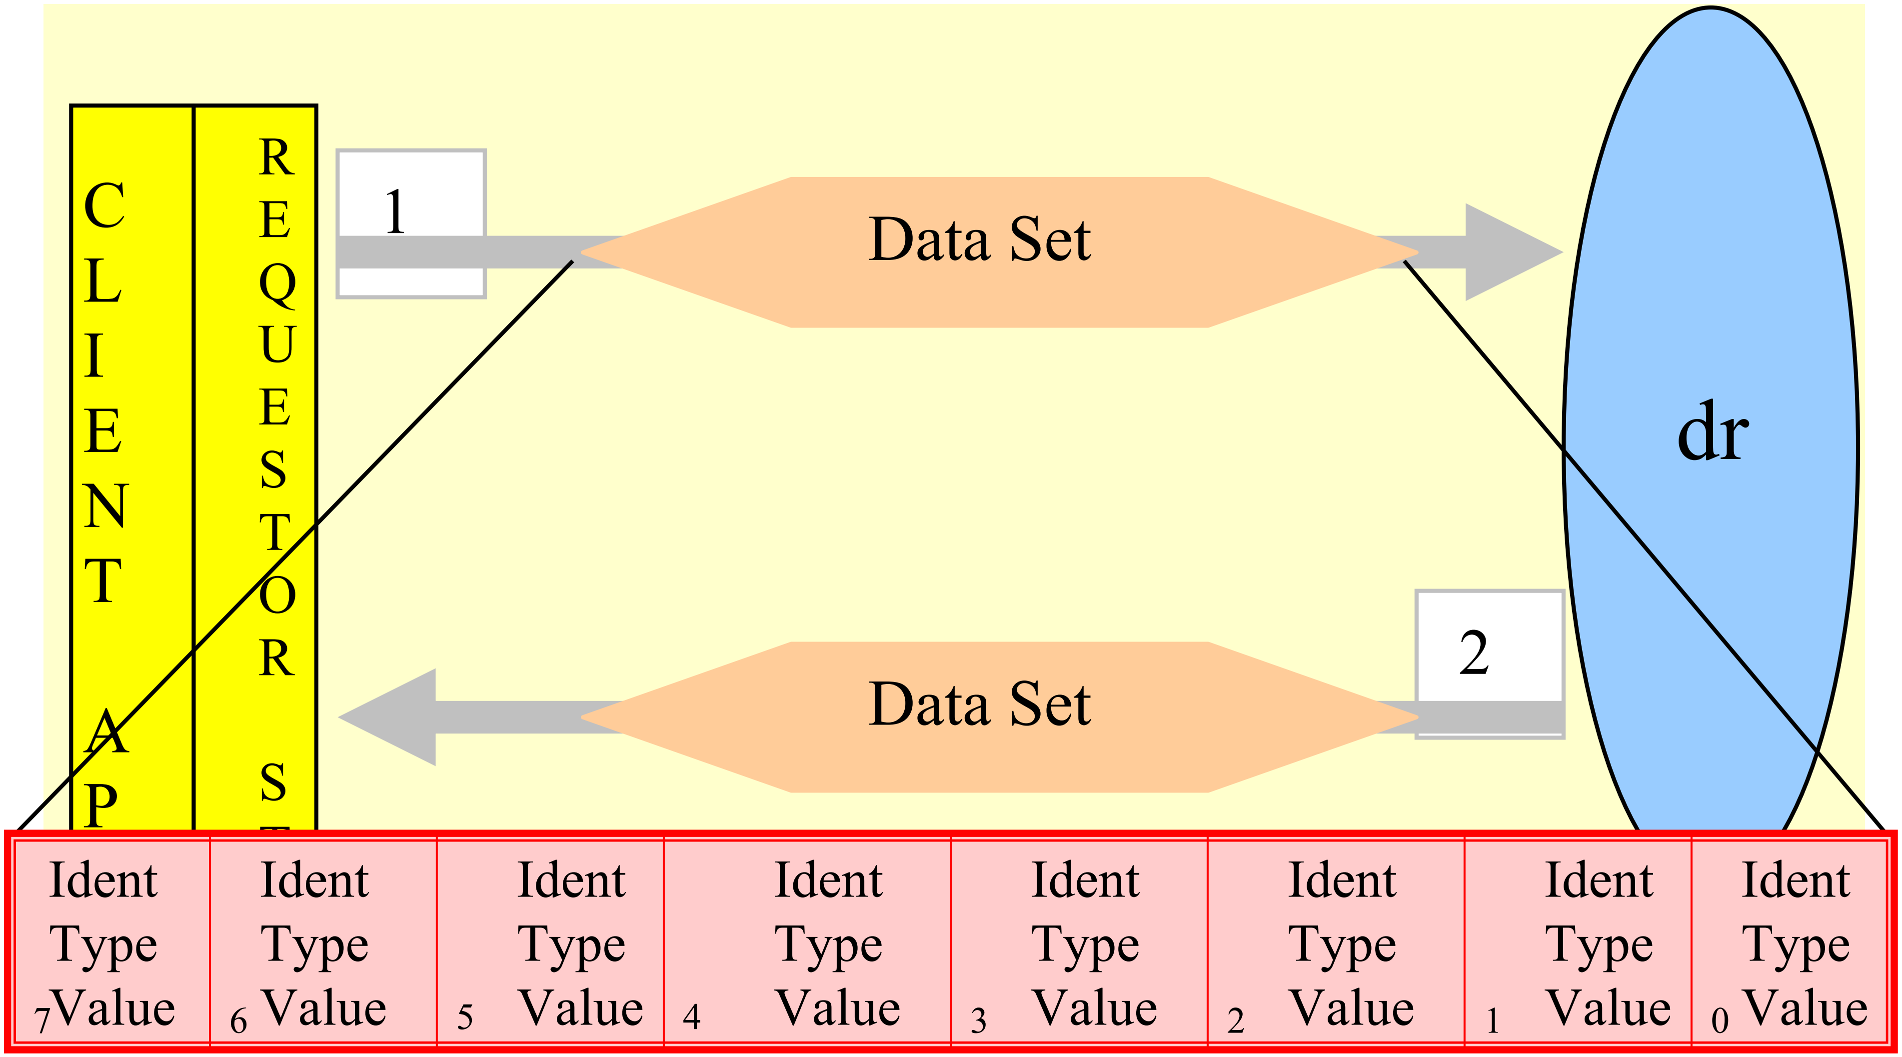

Architecture of Service Interaction: Generic Tasks

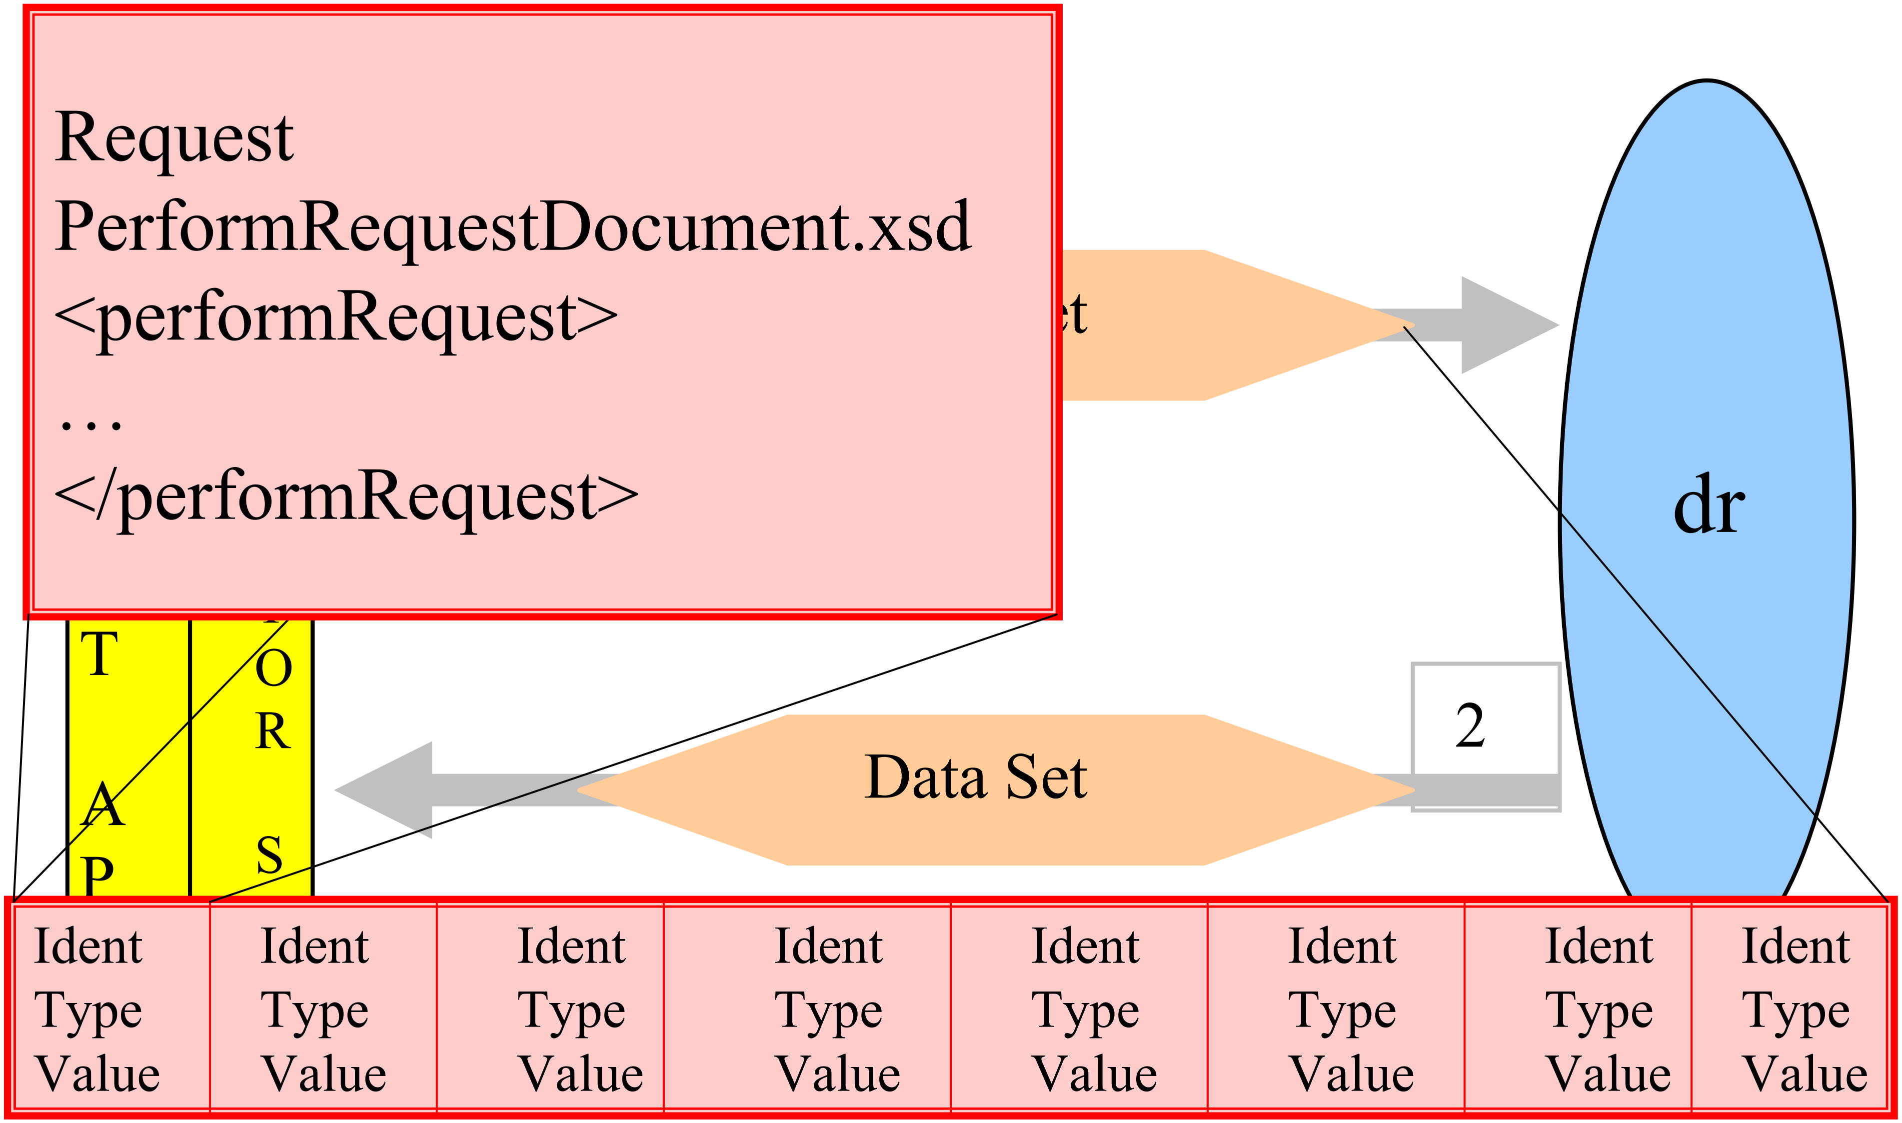

Be Direct

- **Double Handling costs too much**
  - Memory cycles, bus capacity, cache (
- **Double Handling via discs pathol**
- **Data translation expensive**
  - Avoid or compose
- **Main memory is not big enough**
  - Nor is it linear and uniform
  - Streaming algorithms essential
- **Couple generator & consumer di**
  - Data pipe from RAM to RAM
  - Requires coupled computation execu
  - Requires new standards and technol

Breaks down boundaries and merges data, execution & transport requirements.

Demands smart workflow enactment service & foundation services

Future DAI requires Fundamental CS

Enabling Grids for E-science in Europe

- **What Architecture Enables Integration of Data & Computation?**
  - **Common Conceptual Models**
  - **Common Planning & Optimisation**
  - **Common Enactment of Workflows**
  - **Common Debugging**
  - **…**
- **What Fundamental CS is needed?**
  - **Trustworthy code & Trustworthy evaluators**
  - **Decomposition and Recomposition of Applications**
  - **…**
- **Is there an evolutionary path?**
- **Are web services a distraction?**

Summary & Conclusions

Take Home Message

Enabling Grids for E-science in Europe

- **There are plenty of Research Challenges**
  - **Workflow & DB integration, co-optimised**
  - **Distributed Queries on a global scale**
  - **Heterogeneity on a global scale**
  - **Dynamic variability**
    - **Authorisation, Resources, Data & Schema**
    - **Performance**
  - ***Some* Massive Data**
  - **Metadata for discovery, automation, repetition, …**
  - **Provenance tracking**
- **Grasp the theoretical & practical challenges**
  - **Working in Open & Dynamic systems**
  - **Incorporate all computation**
  - **Welcome "code" visiting your data**

Take Home Message (2)

- **Information Grids**
  - **Support for collaboration**
  - **Support for computation and data grids**
  - **Structured data fundamental**
    - **Relations, XML, semi-structured, files, …**
  - **Integrated strategies & technologies needed**
- **OGSA-DAI is here now**
  - **A first step**
  - **Try it**
  - **Tell us what is needed to make it better**
  - **Join in making better DAI services & standards**

Comments & Questions Please

www.ogsadai.org.uk

www.nesc.ac.uk