GGF International Summer School on Grid Computing

Vico Equense (Naples), Italy

Introduction to OGSA-DAI

Prof. Malcolm Atkinson
Director

www.nesc.ac.uk

21st July 2003

Workshop Overview

OGSA-DAI Workshop

- 08:30 – Information Grids & Introduction: Malcolm Atkinson
  - Grids and Virtual Organisations
  - Overview of the architecture
  - Typical end-to-end interaction involving configuration and perform documents – preamble to end-to-end demonstrator: Amy Krause
- 10:30 – Coffee break
- 11:00 – OGSA-DAI Architecture and Configuration: Amy Krause
- 12:15 Lab Session (installation and configuration)
- 13:00 – LUNCH
- 14:00 – Internal Structures of OGSA-DAI: Tom Sugden
  - Low-level architecture
  - Implementing Activities
  - Writing Perform Documents
- 15:00 – Lab session (configuration and perform documents)
- 16:30 – BREAK
- 17:00 – Lab Session (Writing your own perform documents)
  - Playtime with OGSA-DAI
- 18:30 – End of Lab sessions

Outline

- What is e-Science?
  - Grids, Collaboration, Virtual Organisations
  - Structured Data at its Foundation
- Motivation for DAI
  - Key Uses of Distributed Data Resources
  - Challenges
- Introduction to DAI
  - GGF DAIS Working Group
  - Conceptual Models
  - Architectures
  - Current OGSA-DAI components

e-Science & Grids

It's Easy to Forget How Different 2003 is From 1993

- Enormous quantities of data: Petabytes
  - For an increasing number of communities, gating step is not collection but analysis
- Ubiquitous Internet: >100 million hosts
  - Collaboration & resource sharing the norm
  - Security and Trust are crucial issues
- Ultra-high-speed networks: >10 Gb/s
  - Global optical networks
  - Bottlenecks: last kilometre & firewalls
- Huge quantities of computing: >100 Top/s
  - Moore's law gives us all supercomputers
  - Ubiquitous computing
- Moore's law everywhere
  - Instruments, detectors, sensors, scanners, …

Derived from Ian Foster's slide at ssdbM July 03

Foundation for e-Science

- e-Science methodologies will **rapidly transform** science, engineering, medicine and business
  - driven by exponential growth (×1000/decade)
    - enabling a whole-system approach

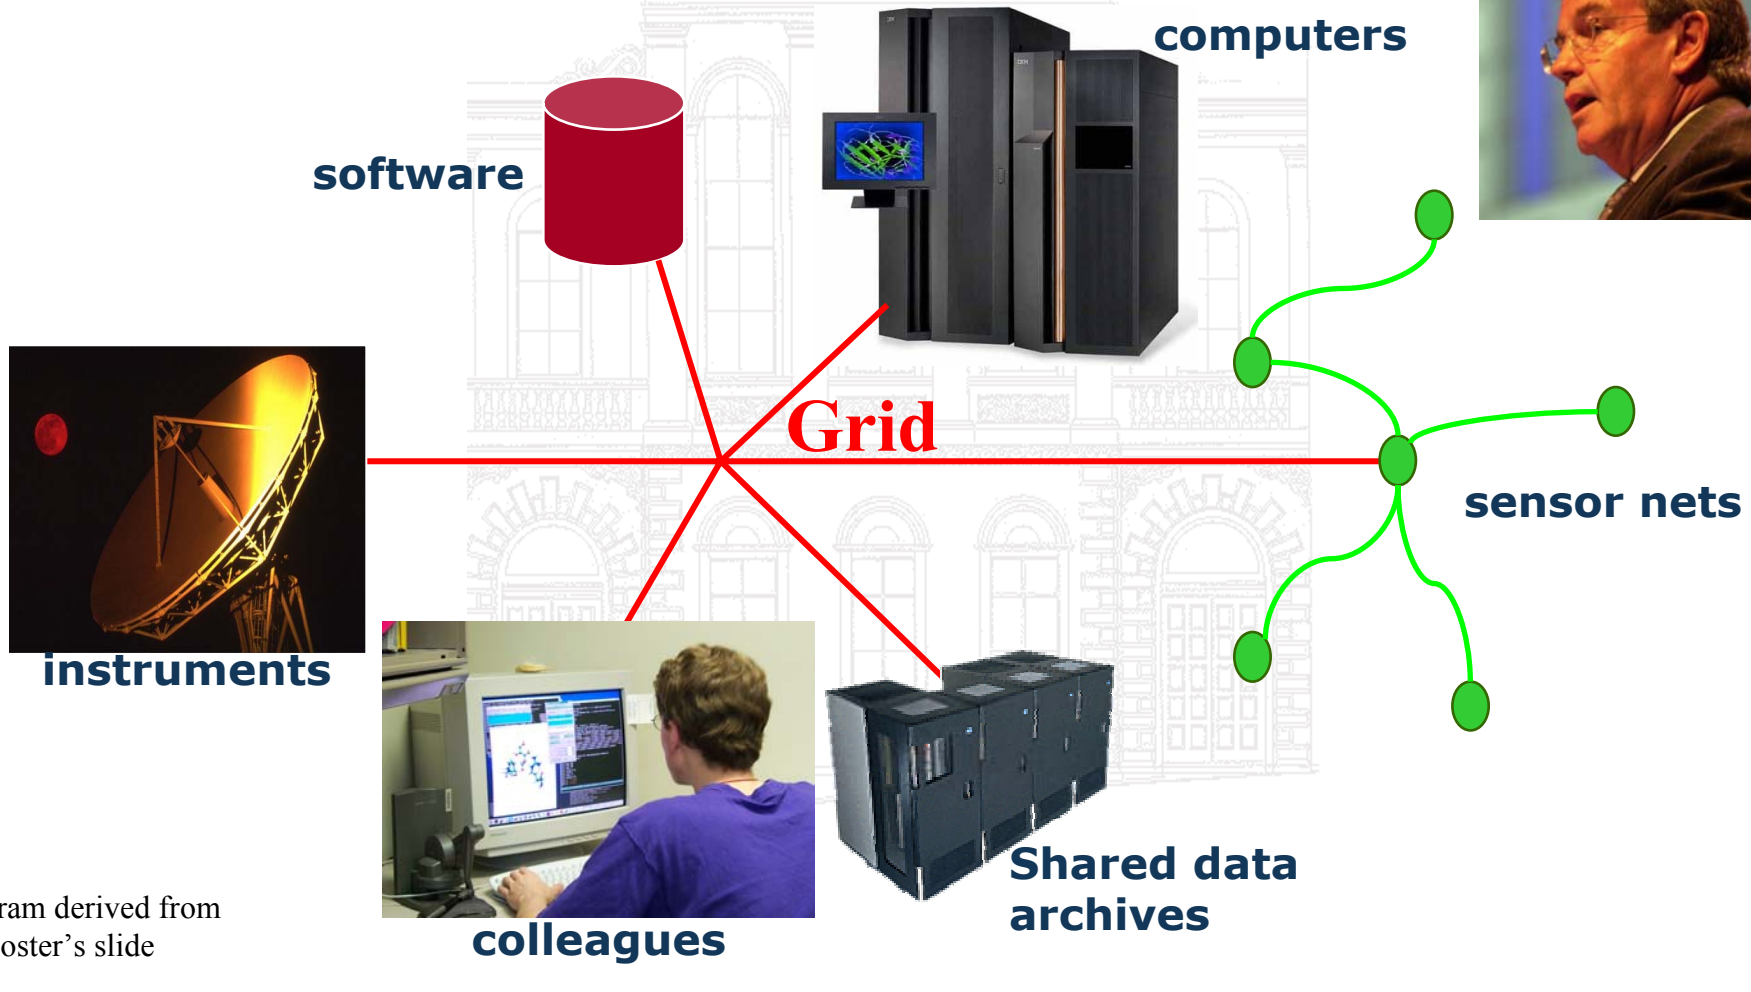

Diagram derived from Ian Foster's slide

e-Science & Collaboration

Three-way Alliance

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

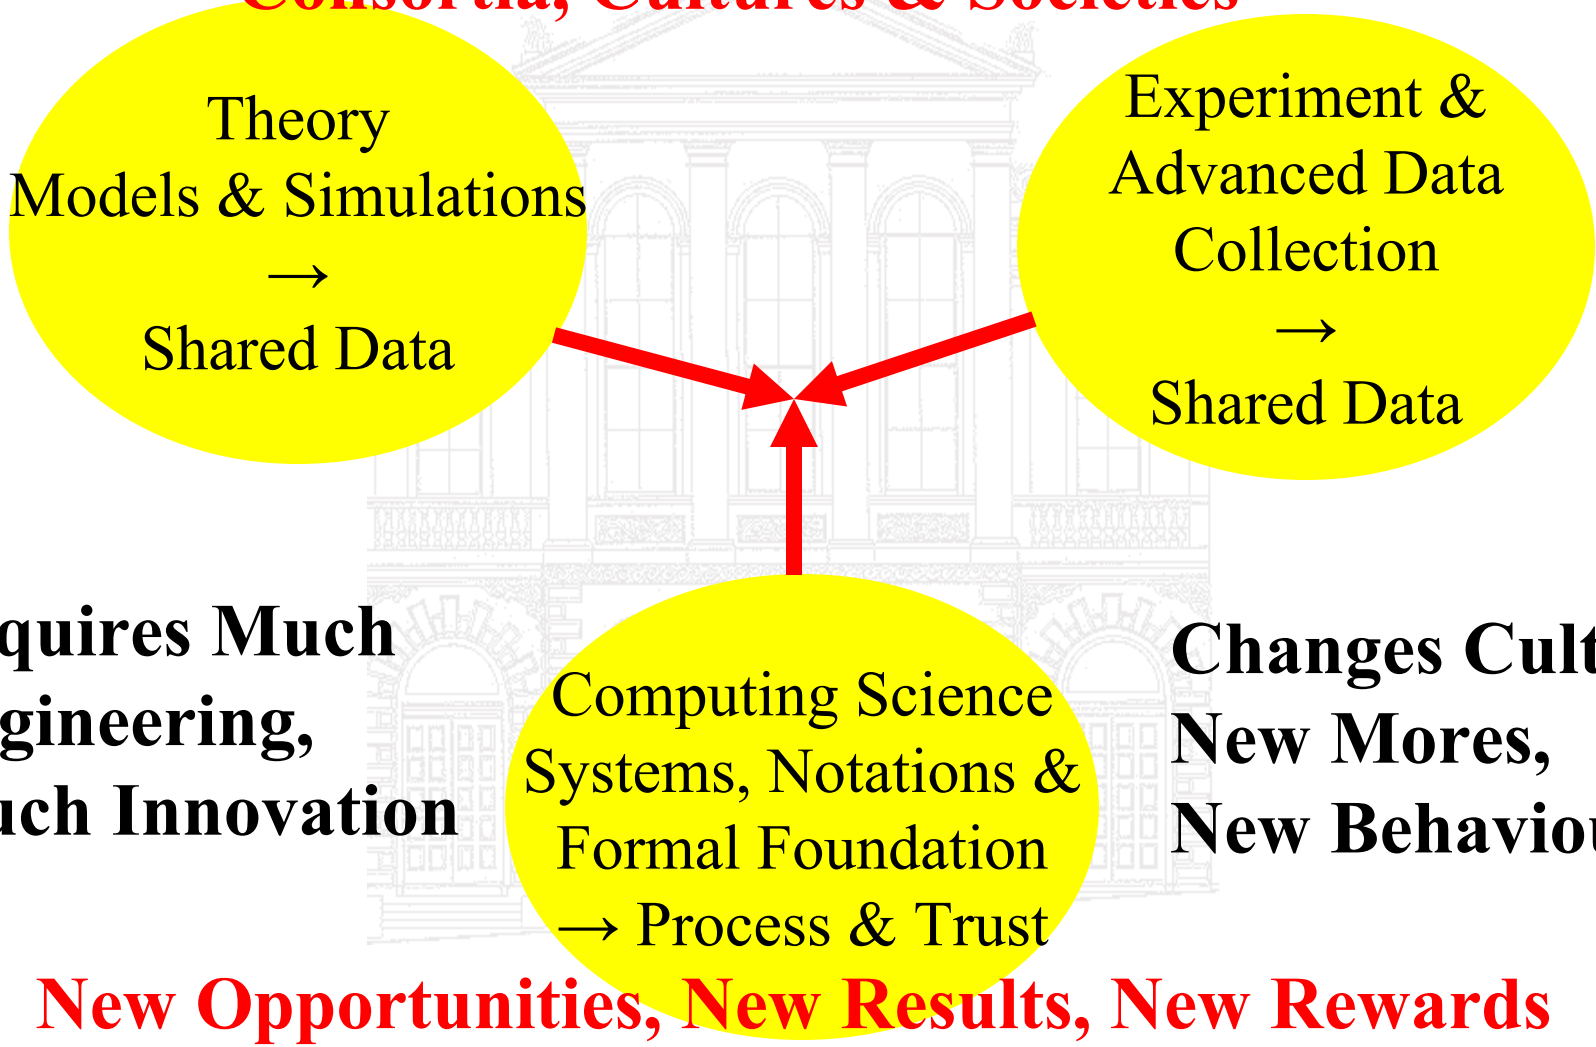

Requires Much Engineering, Much Innovation

Changes Culture, New Mores, New Behaviours

New Opportunities, New Results, New Rewards

Biochemical Pathway Simulator

(Computing Science, Bioinformatics, Beatson Cancer Research Labs)

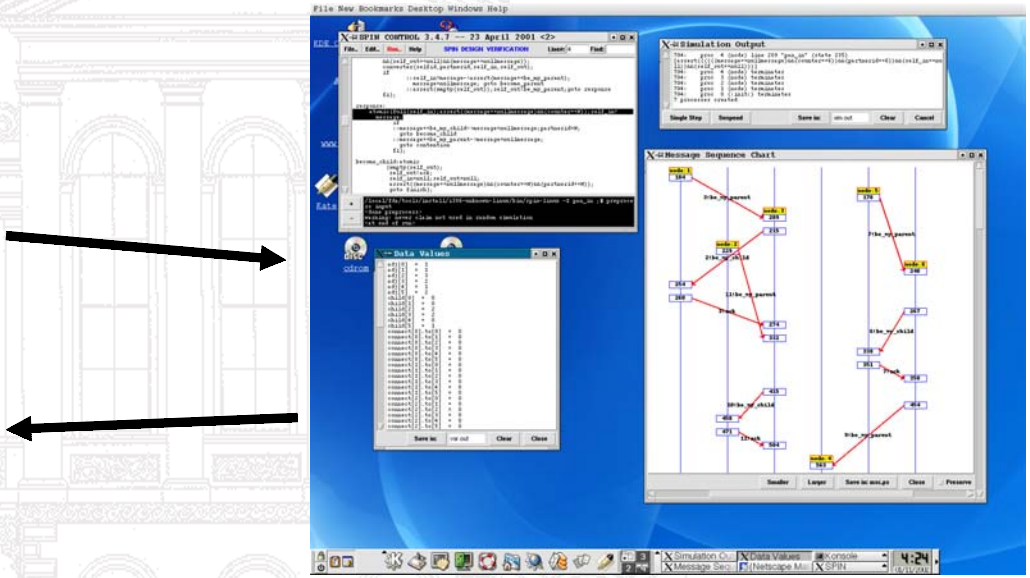

Closing the information loop – between lab and computational model.

DTI Bioscience *Beacon Project*

Harnessing Genomics Programme

Slide from Muffy Calder, Glasgow

e-Science, Virtual Organisations & Knowledge Communities

Emergence of Global Knowledge Communities

- Teams organised around common goals
  - Communities: "Virtual organisations"
  - Overlapping memberships, resources and activities
- Essential diversity is a strength & challenge
  - membership & capabilities
- Geographic and political distribution
  - No location/organisation/country possesses all required skills and resources
- Dynamic: adapt as a function of their situation
  - Adjust membership, reallocate responsibilities, renegotiate resources

Slide derived from Ian Foster's ssdbm 03 keynote

The Emergence of Global Knowledge Communities

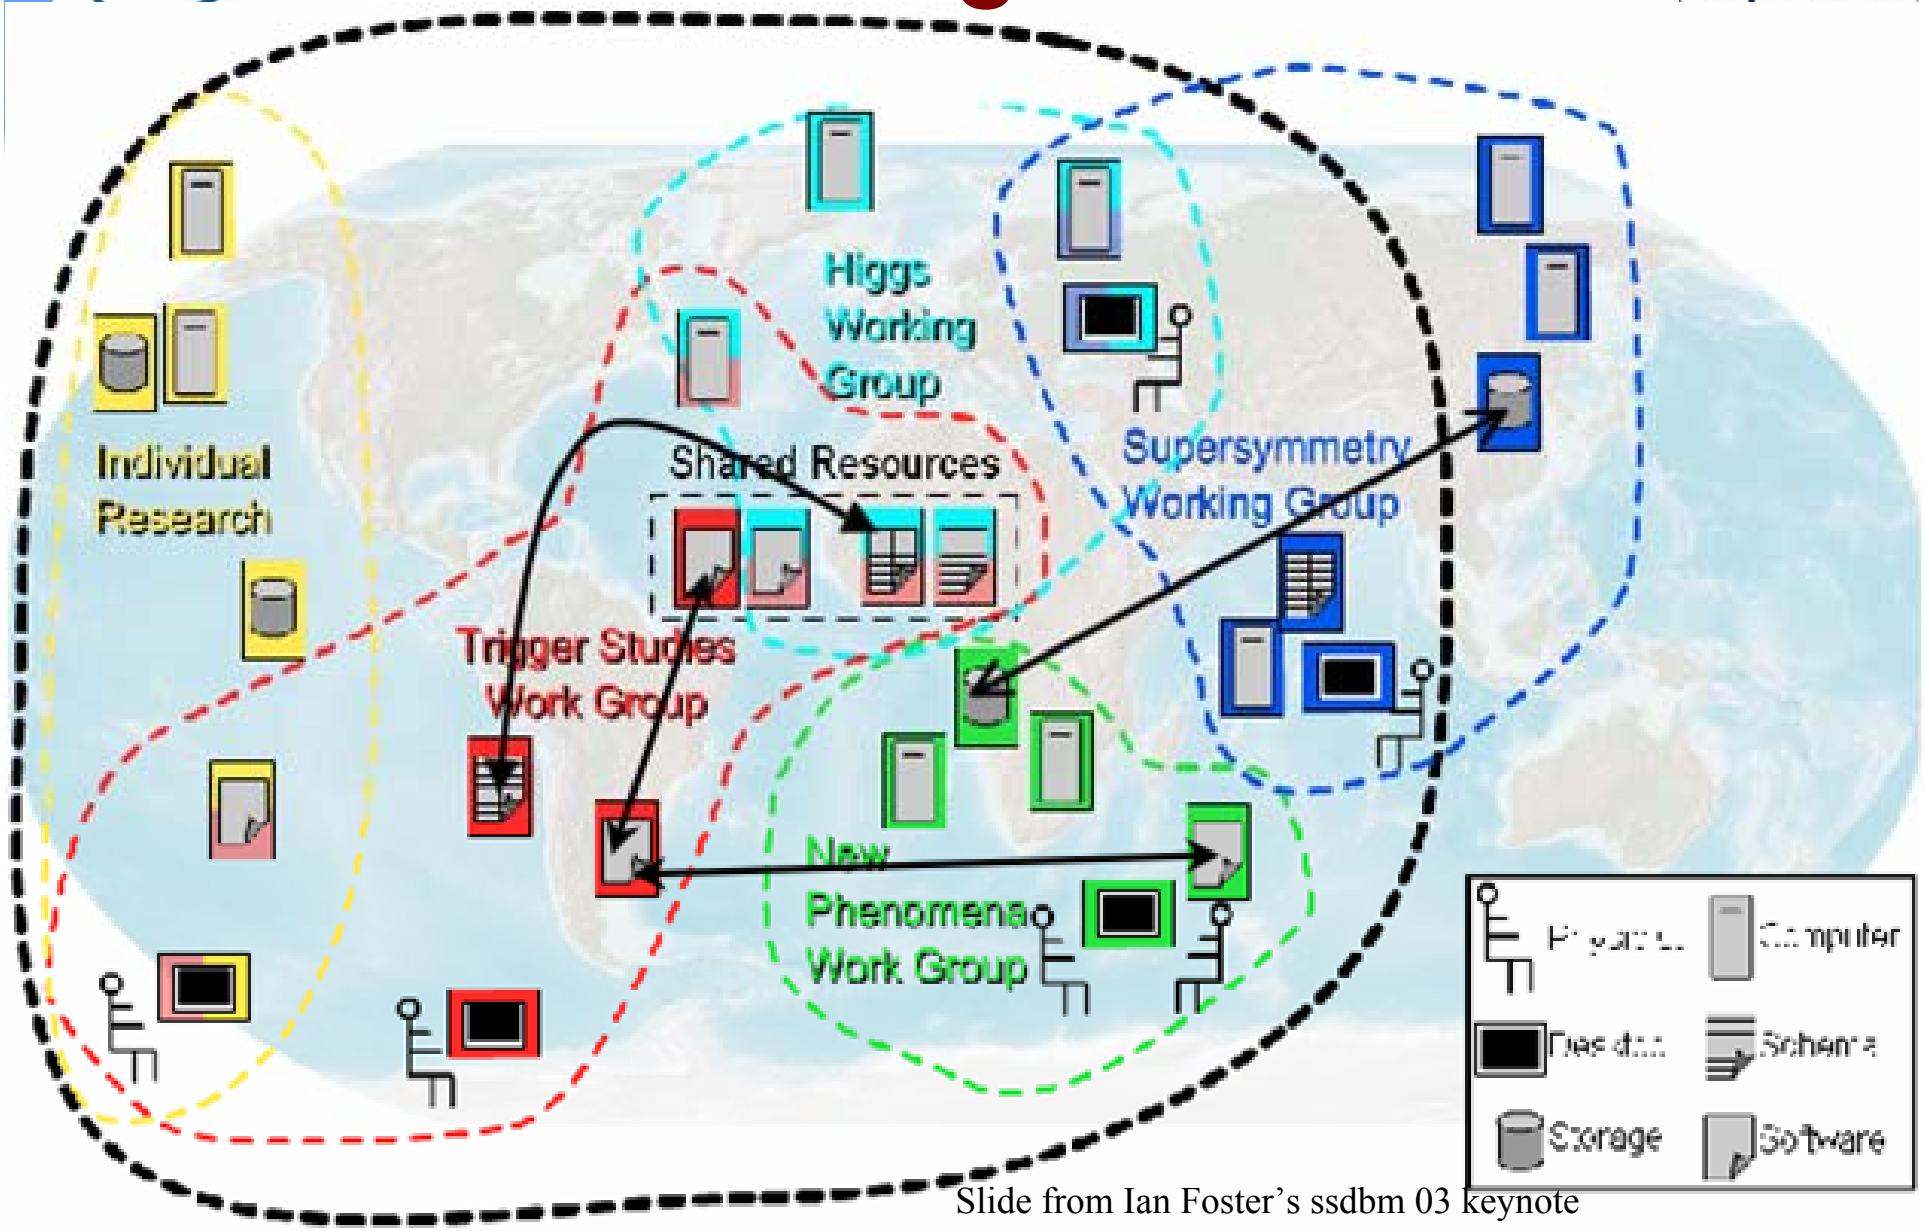

Slide from Ian Foster's ssdbm 03 keynote

Global Knowledge Communities Often Driven by Data: E.g., Astronomy

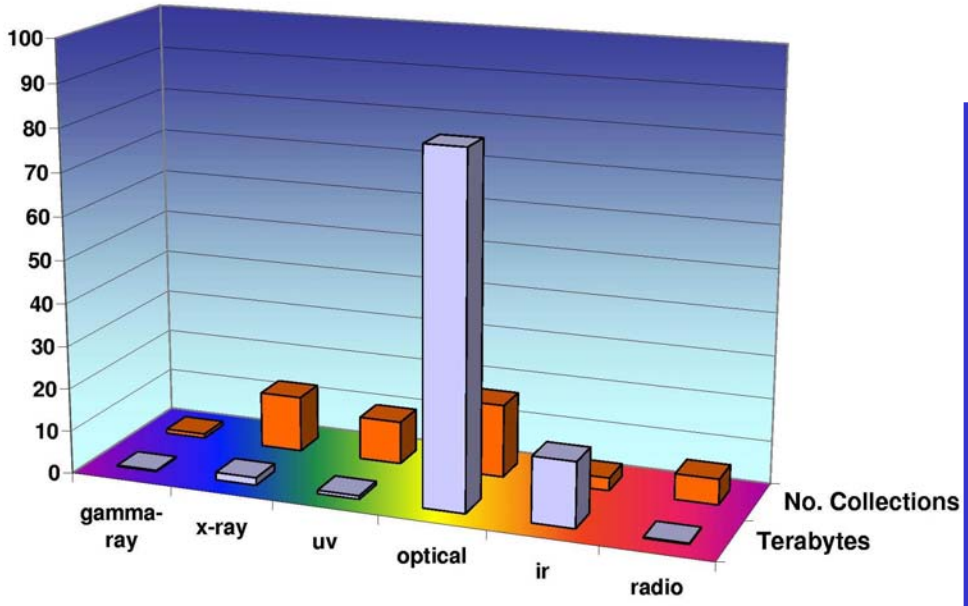

No. & sizes of data sets as of mid-2002, grouped by wavelength

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

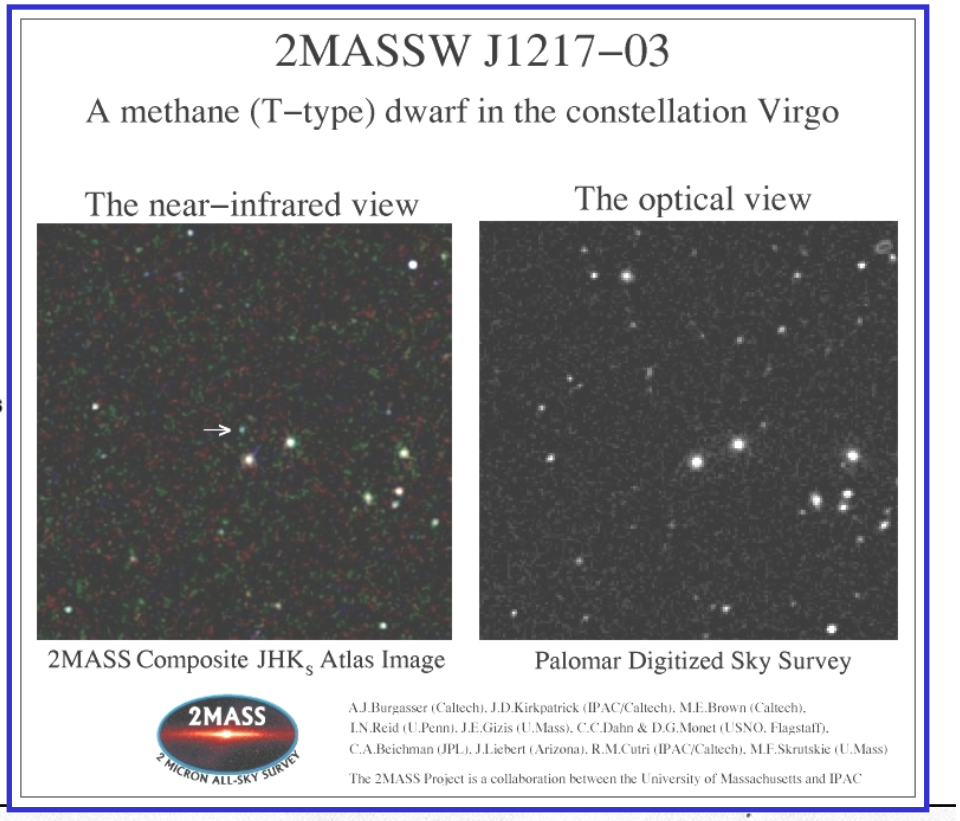


THE NEW YORK TIMES **NATIONAL** TUESDAY, JUNE 1, 1999

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

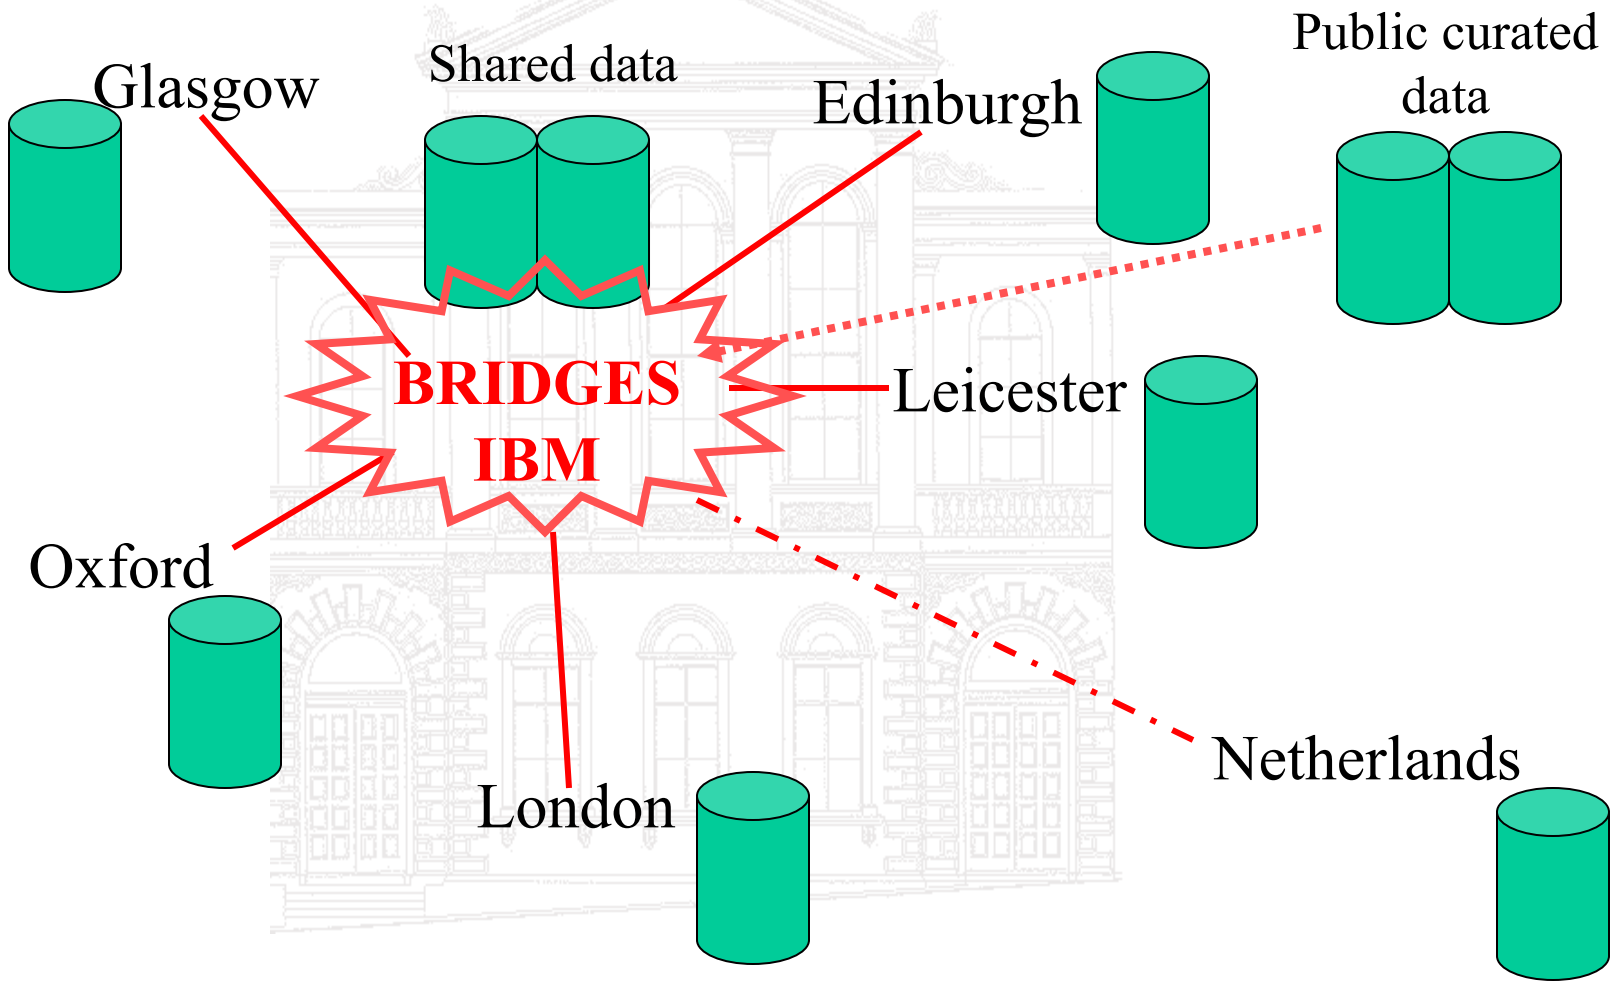

Database-mediated Communication

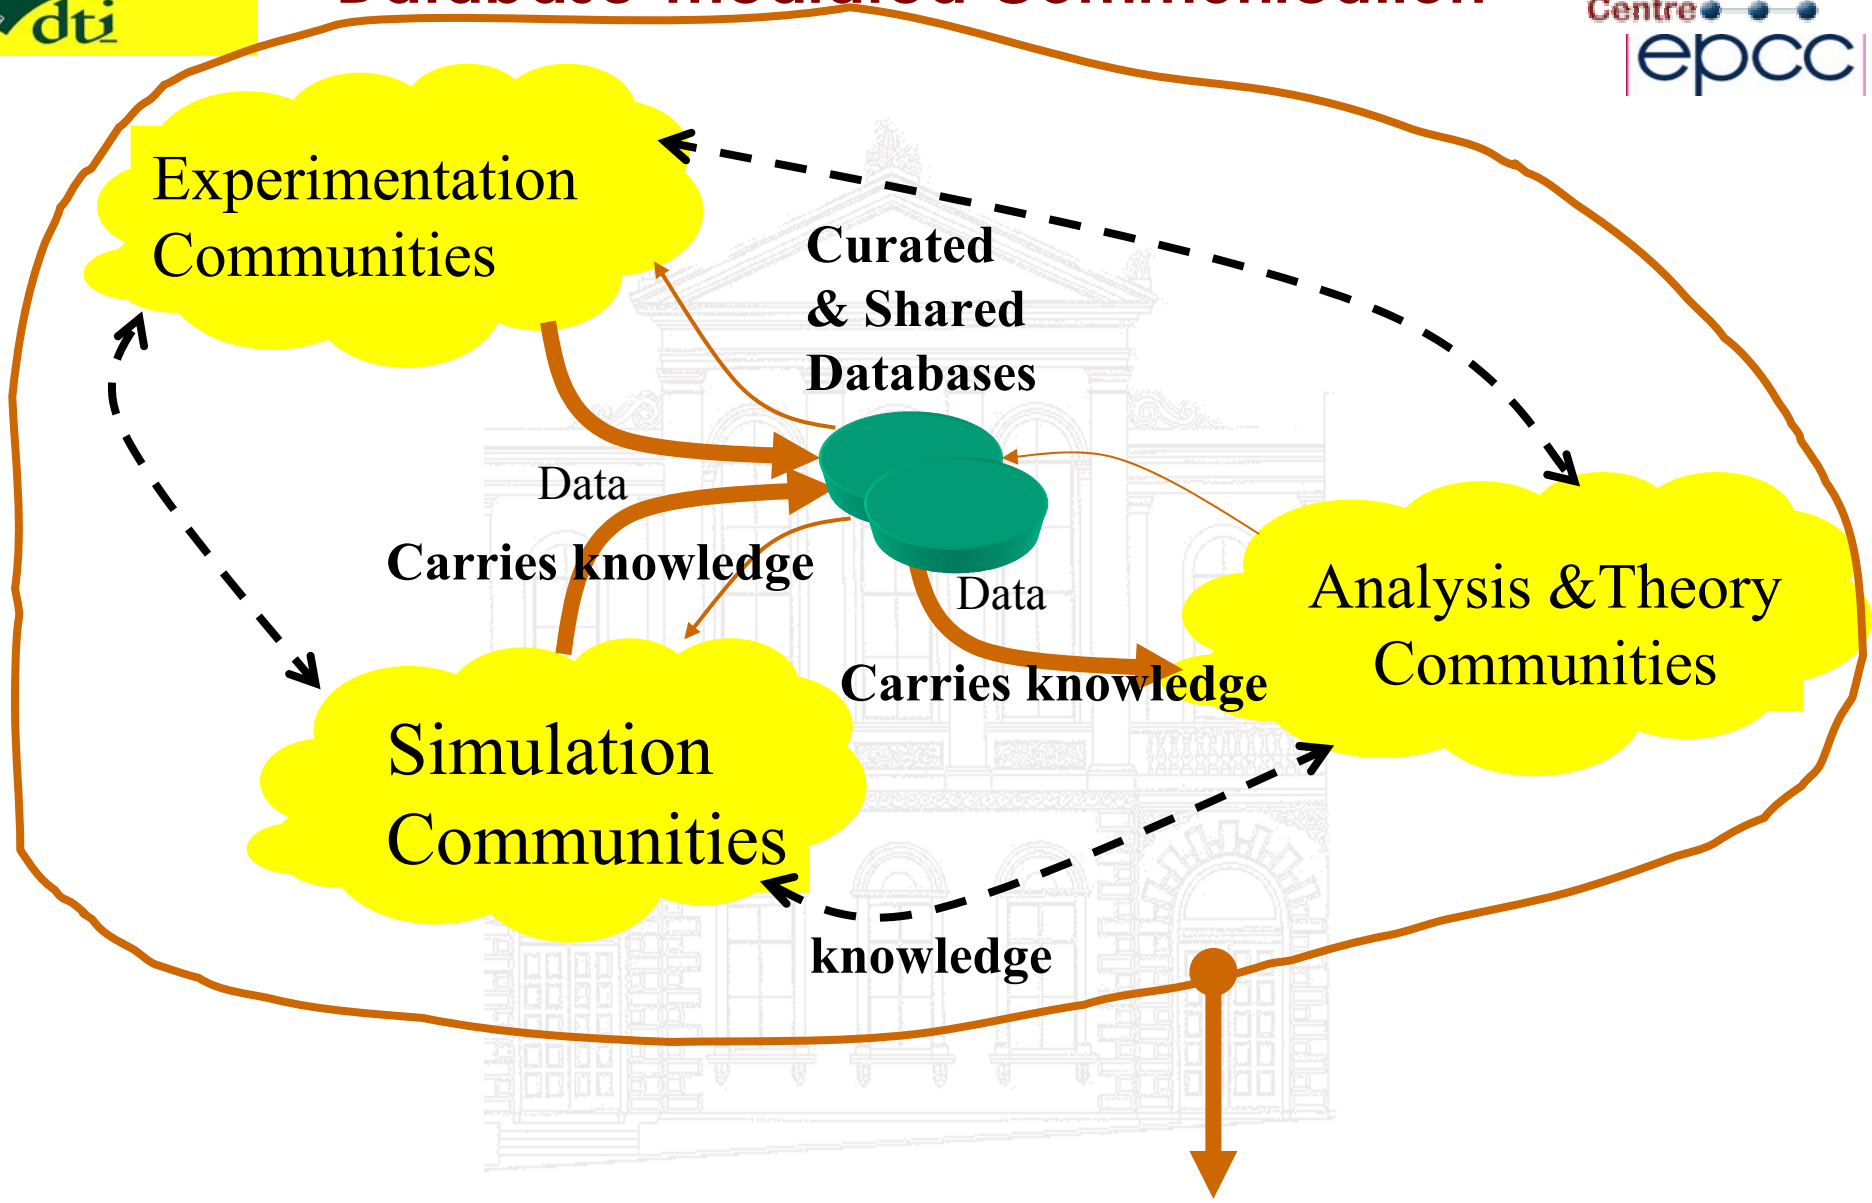

e-Science, Data Scales, Challenges & Opportunities

global in-flight engine diagnostics

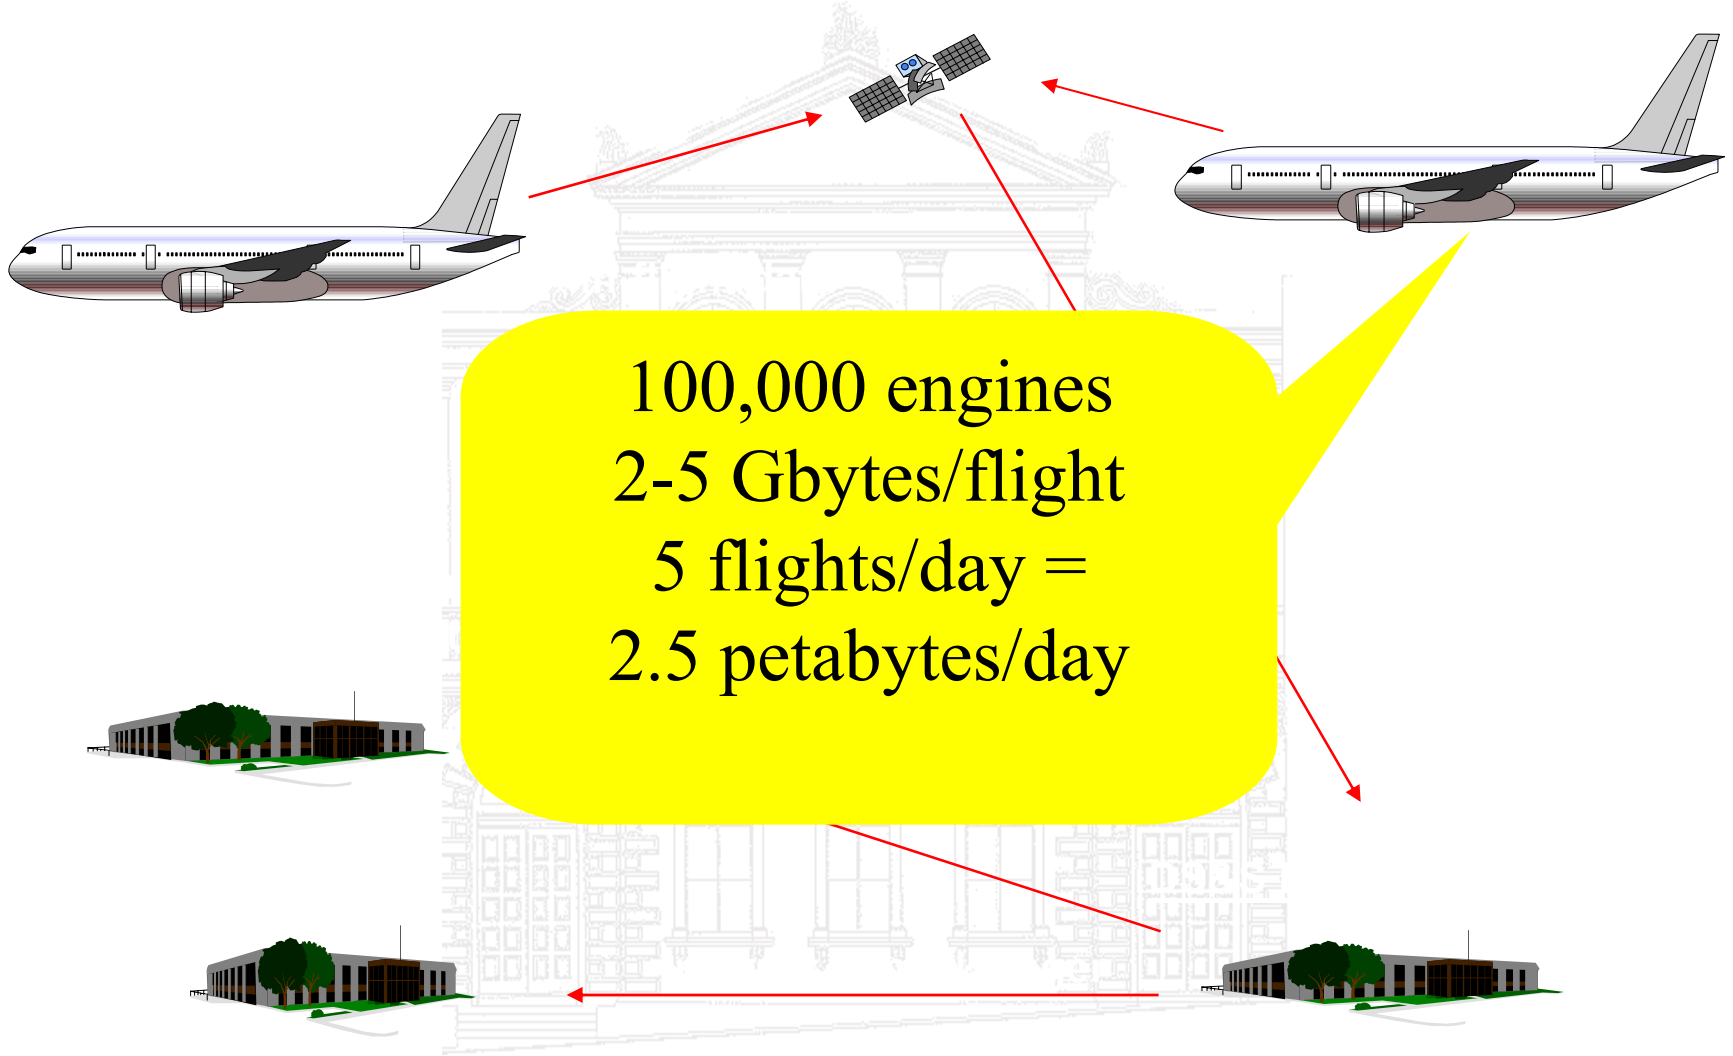

Database Growth

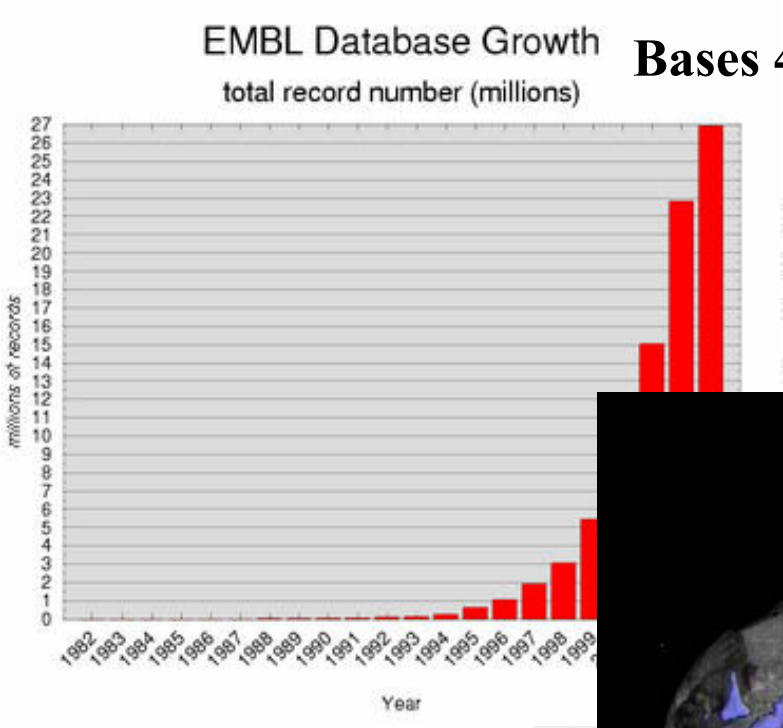

Bases 41,073,690,490

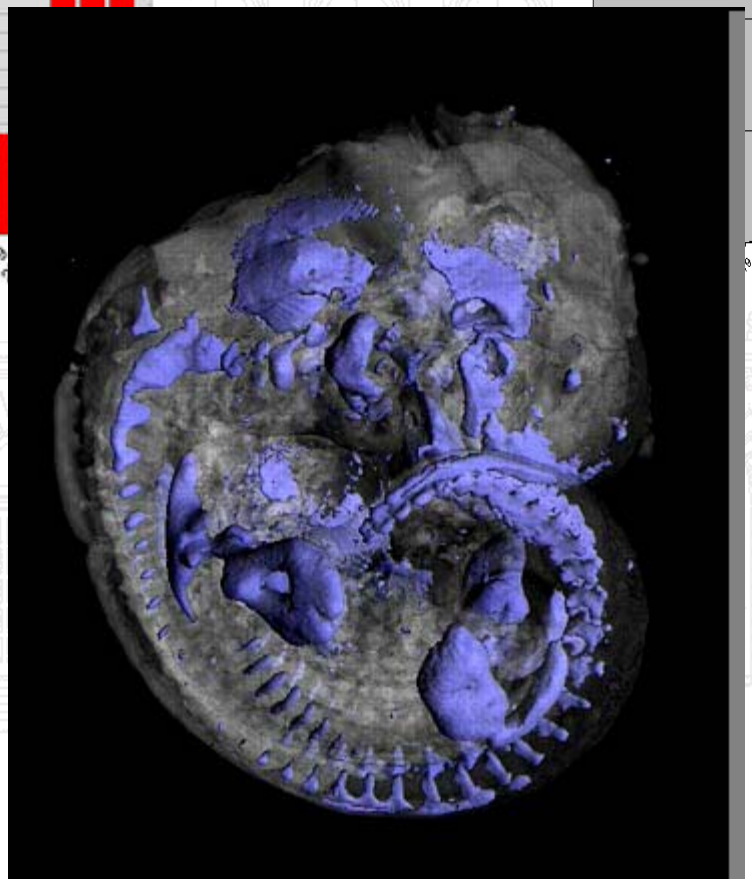

PDB Content Growth

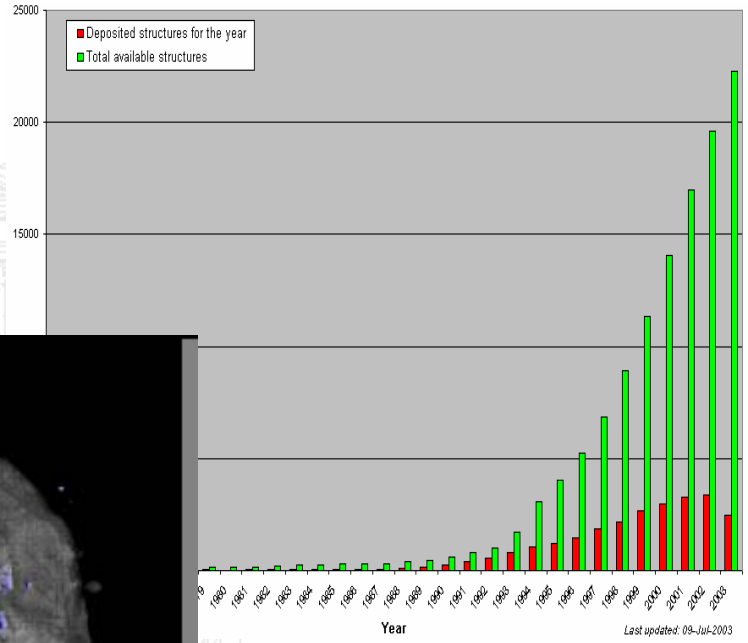

Distributed Structured Data

- Key to Integration of Scientific Methods
- Key to Large-scale Collaboration
- Many Data Resources
  - Independently managed
  - Geographically distributed
  - Primary Data, Data Products, Meta Data, Administrative data, …
- Discovery and Decisions!
  - Extracting nuggets from multiple sources
  - Combing them using sophisticated models
  - Analysis on scales required by statistics
- Repeated Processes

Petabyte of Digital Data / Hospital / Year

Tera → Peta Bytes

- RAM time to move
  - 15 minutes
- 1Gb WAN move time
  - 10 hours ($1000)
- Disk Cost
  - [illegible] Disks + 490 units + 32 racks = $7 million
- Disk Power
  - 100 Kilowatts
- Disk Weight
  - 5.6 Kg
  - 33 Tonnes
- Disk Footprint
  - Inside machine
  - 60 m²

- RAM time to move
  - 2 months
- 1Gb WAN move time
  - 14 months

How make it secure & reliable!

May 2003 Approximately Correct

See also *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?**

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 the sensor & simulation architectures for this?
- 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

Scientific Data

- **Opportunities**
  - Global Production of *Published* Data
  - Volume↑ Diversity↑
  - Combination ⇒ Analysis ⇒ Discovery
- **Challenges**
  - Data Huggers
  - Meagre metadata
  - Ease of Use
  - Optimised integration
  - Dependability
- **Opportunities**
  - Specialised Indexing
  - 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 & …
- Information Grids will be very important
- 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

Outline

- What is e-Science?
  - Grids, Collaboration, Virtual Organisations
  - Structured Data at its Foundation
- Motivation for DAI
  - Key Uses of Distributed Data Resources
  - Challenges
- **Introduction to Data Access & Integration**
  - **DAIS-WG: Conceptual Model & Architecture**
  - **Data Access & Integration in OGSA**
  - **Introducing OGSA-DAI Services**
- **Looking ahead & Take-Home Messages**
  - **Composition of Analysis & Interpretation**

Science as Workflow

- Data integration = the derivation of new data from old, via coordinated computation
- May be computationally demanding
  - The workflows used to achieve integration are often valuable artifacts in their own right

- **May be Data Access & Movement Demanding**
  - Obtaining data from files and DBs, transfer between computations, deliver to DBs and File stores

- Thus we must be concerned with how we
  - Build workflows
  - Share and reuse workflows
  - Explain workflows
  - Schedule workflows

- **Consider also DBs & (Autonomous) Updates**
- **External actions are important**

Slide derived from Ian Foster's ssdbm 03 keynote

Sloan Digital Sky Survey Production System

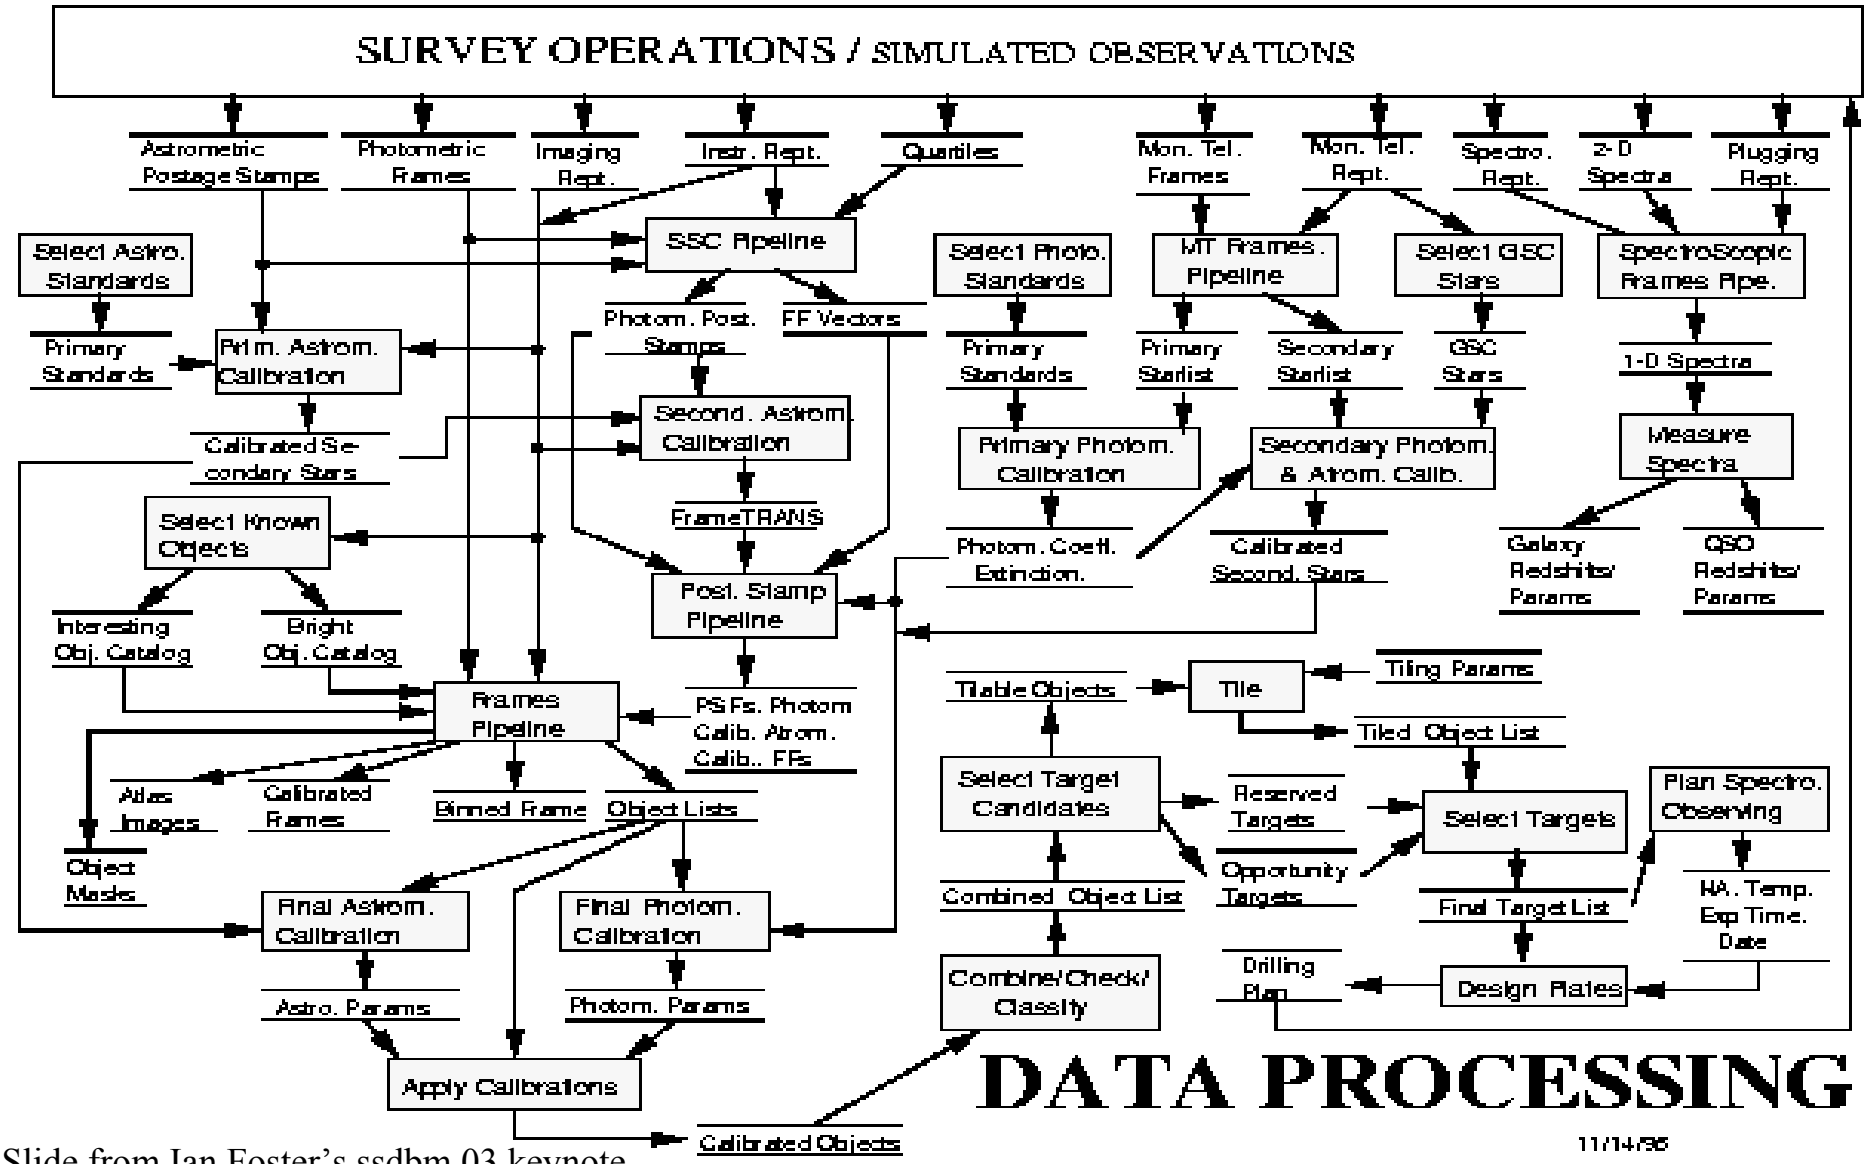

Slide from Ian Foster's ssdbm 03 keynote

DAIS WG

DAIS-WG

- Specification of Grid Data Services
- Chairs
  - Norman Paton, Manchester University
  - Dave Pearson, Oracle
- Current Spec. Draft Authors

Mario Antonioletti
Malcolm Atkinson
Neil P Chue Hong
Amy Krause
Susan Malaika
Gavin McCance
Simon Laws
James Magowan
Norman W Paton
Greg Riccardi

Draft Specification for GGF 7

Grid Database Service Specification

Status of This Memo

This memo provides information to the Grid community regarding the specification of Grid Database Services. The specification is presently a draft for discussion. It does not define any standards or technical recommendations. Distribution is unlimited.

Abstract

Data management systems are central to many applications across multiple domains, and play a significant role in many others. Web services provide implementation neutral facilities for describing, invoking and orchestrating collections of networked resources. The Open Grid Services Architecture (OGSA) extends Web Services with consistent interfaces for creating, managing and exchanging information among Grid Services, which are dynamic computational artefacts cast as Web Services. Both Web and Grid service communities stand to benefit from the provision of consistent, agreed service interfaces to database management systems. Such interfaces must support the description and use of database systems using Web Service standards, taking account of the design conventions and mandatory features of Grid Services. This document presents a specification for a collection of Grid Database Services. The proposal is presented for discussion within the Global Grid Forum (GGF) Database Access and Integration Services (DAIS) Working Group, in the hope that it will evolve into a formal standard for Grid Database Services. There are several respects in which the current proposal is incomplete, but it is hoped that the material included is sufficient to allow an informed discussion to take place concerning both its form and substance.

Conceptual Model
External Universe

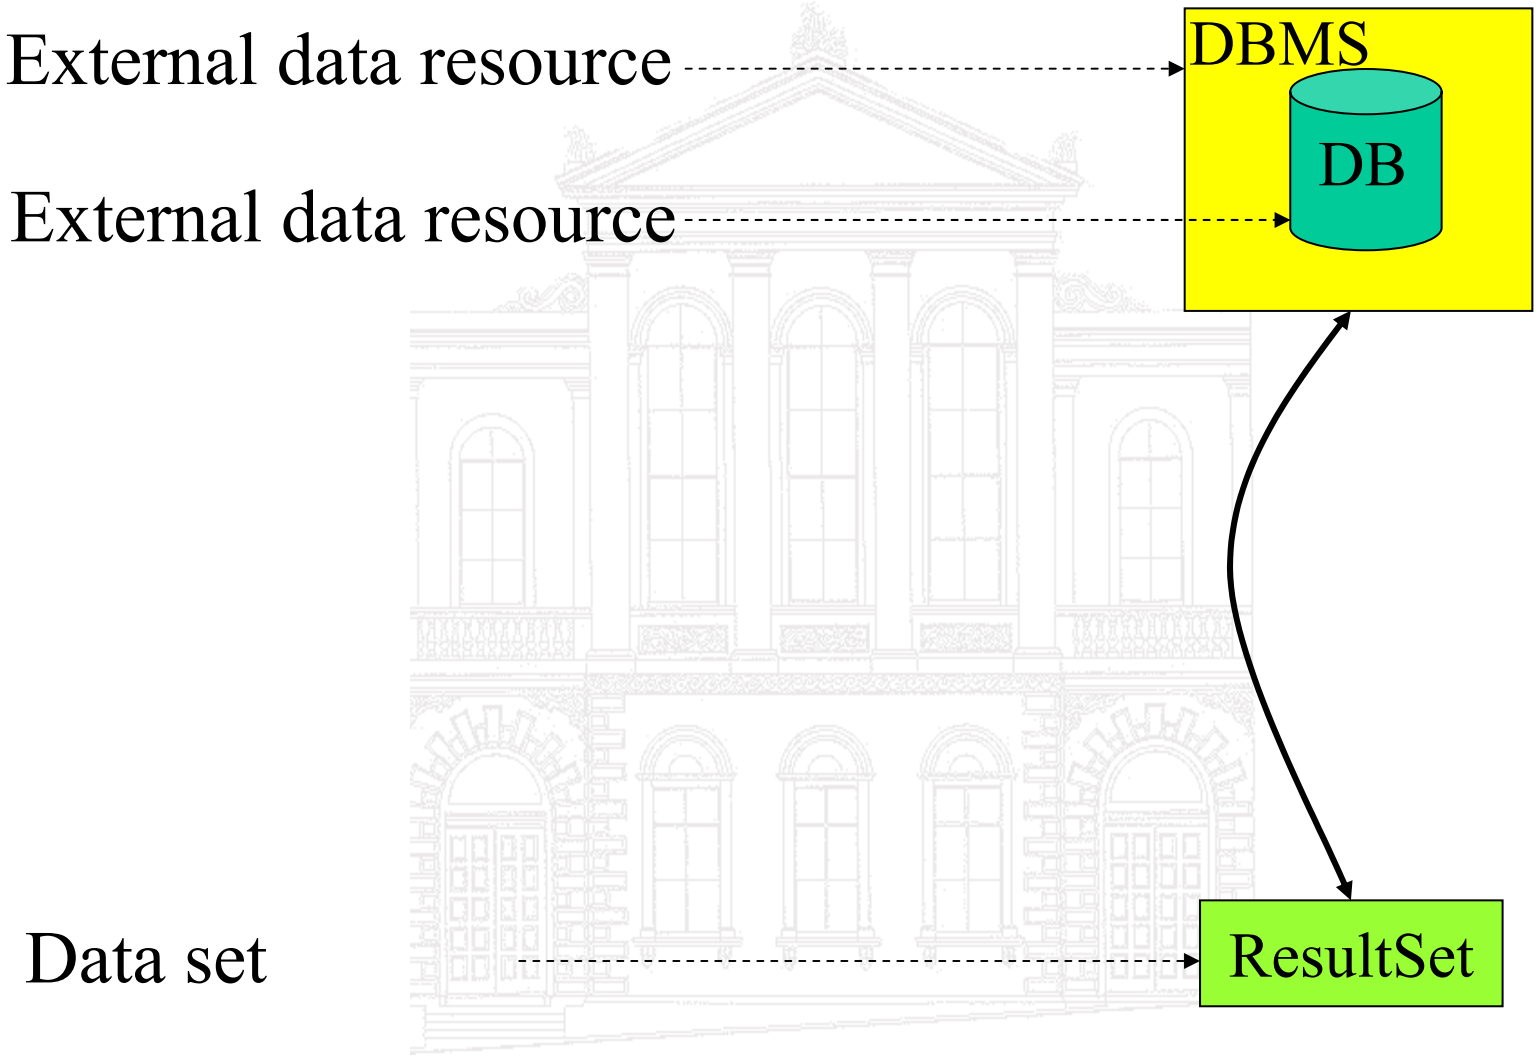

Conceptual Model
DAI Service Classes

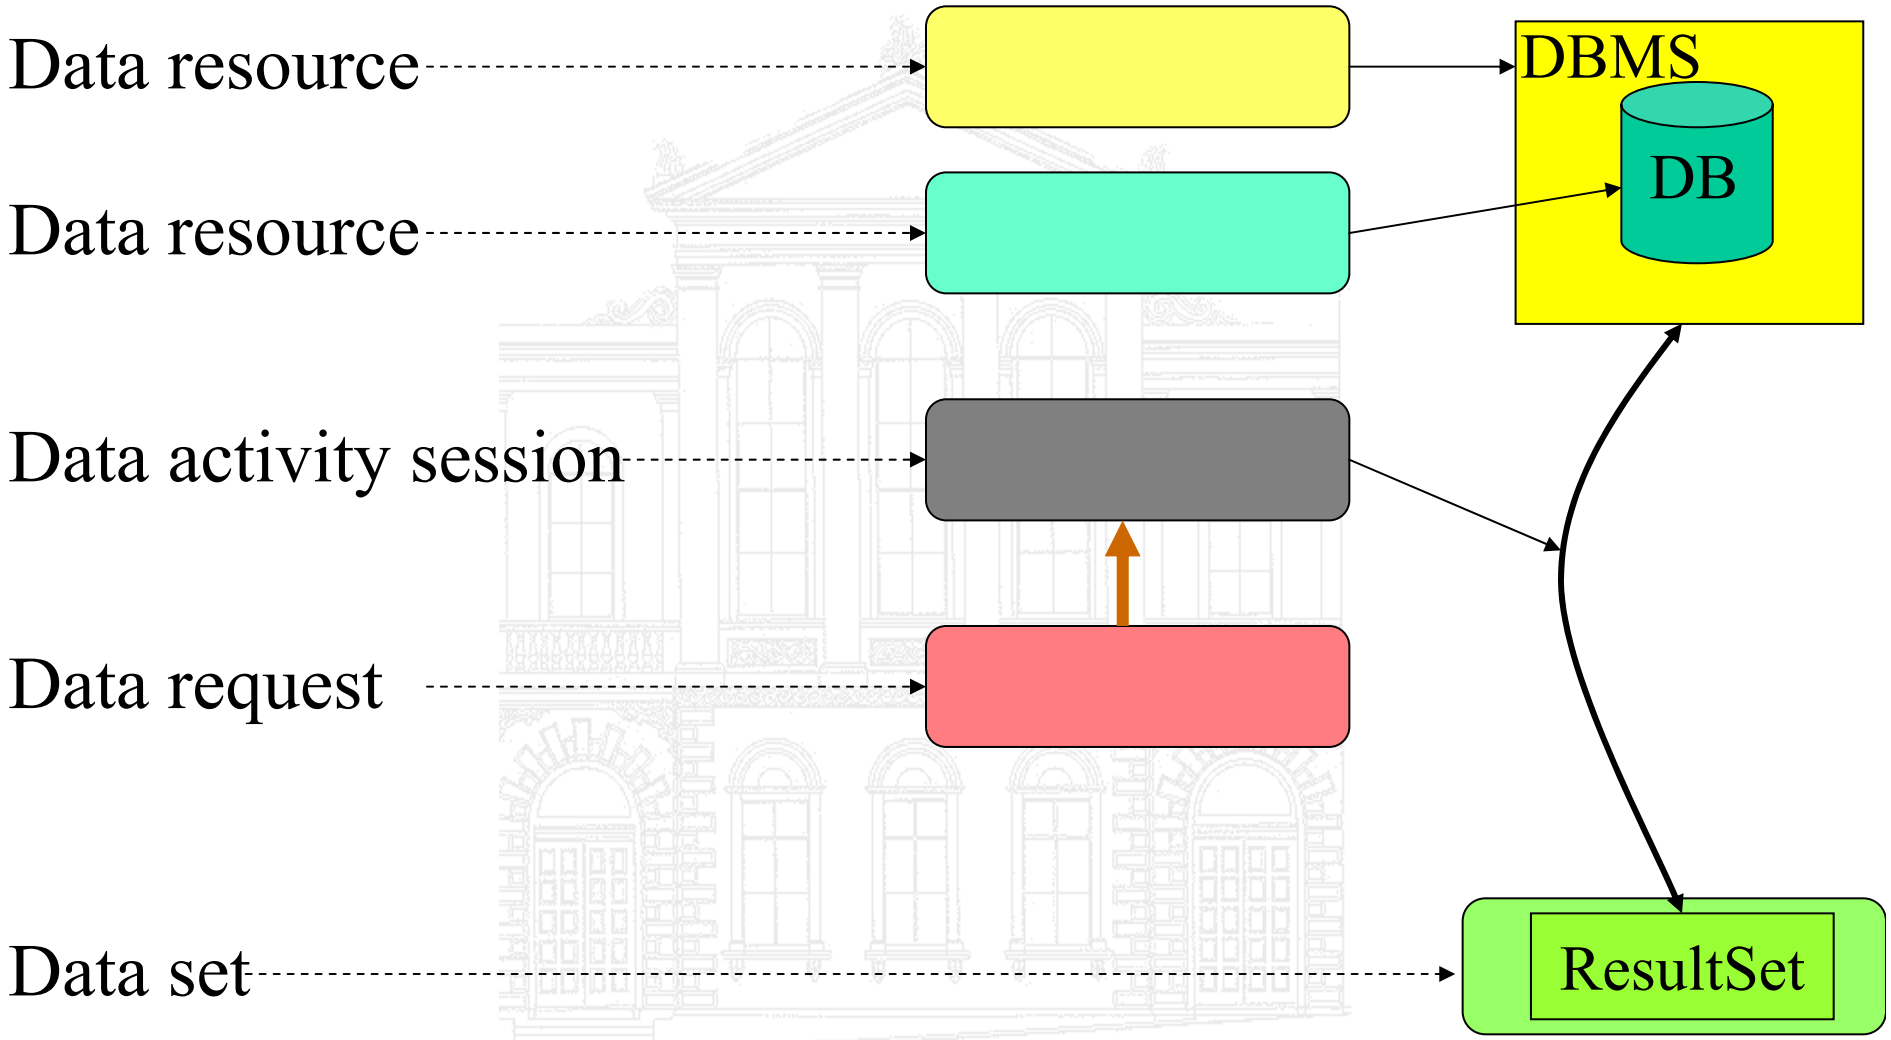

Architecture of Service Interaction

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

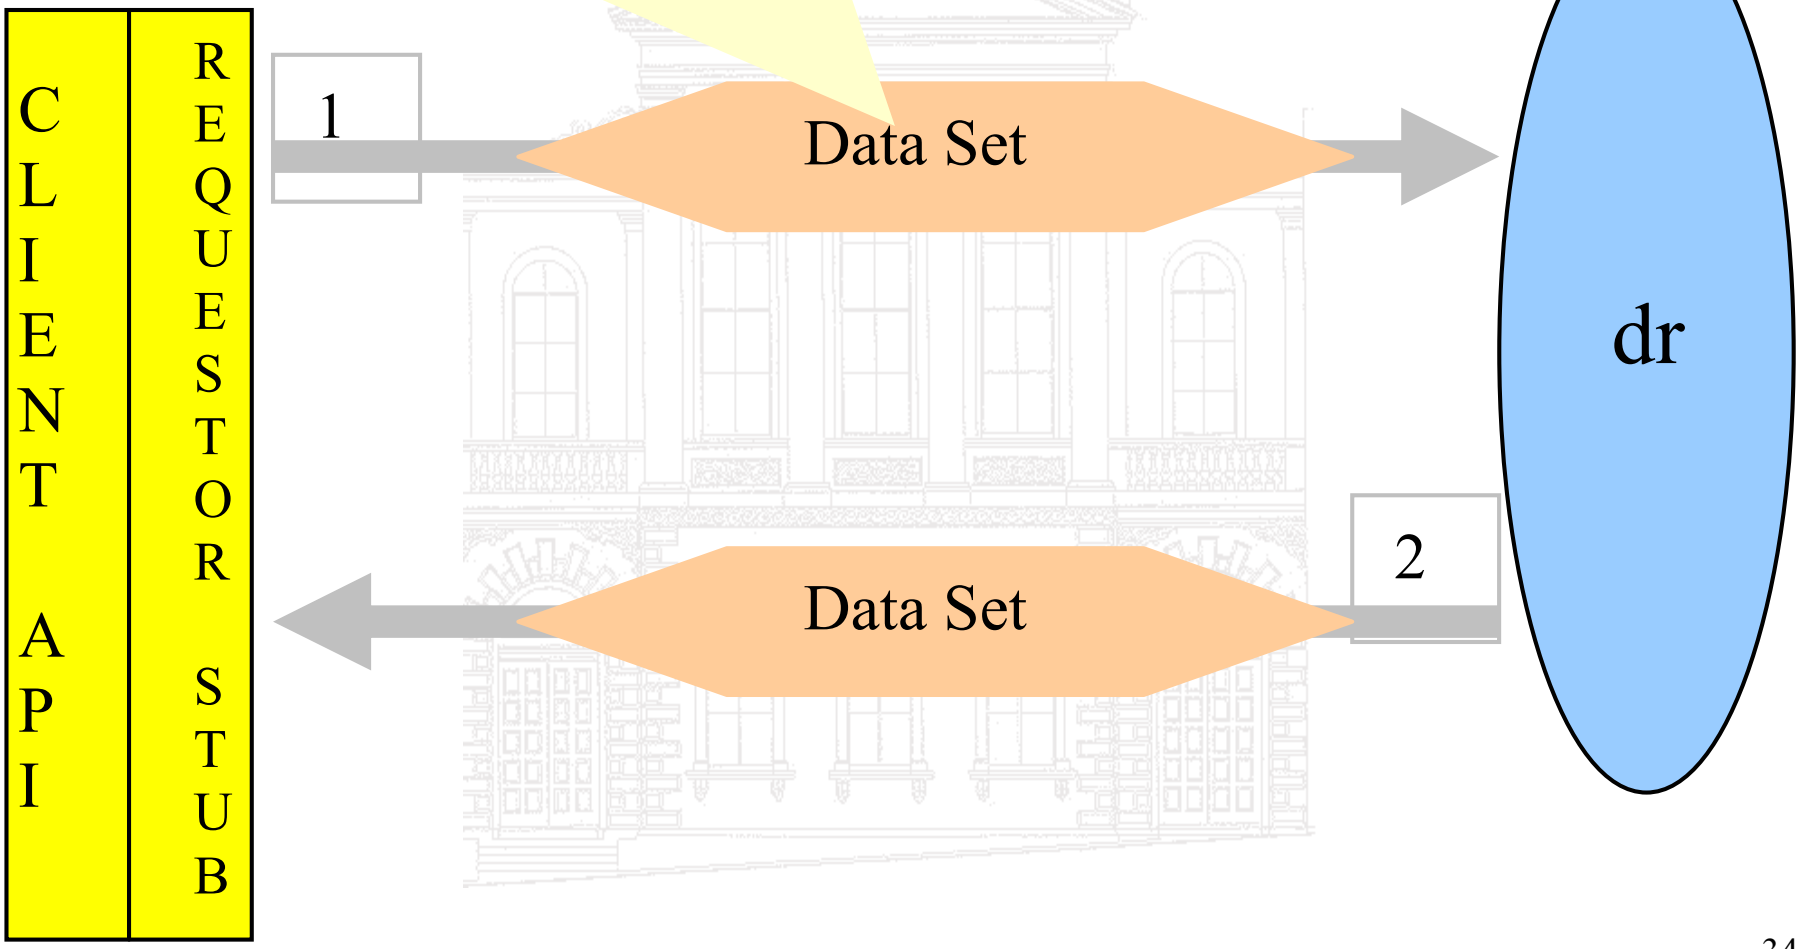

Architecture of Service Interaction

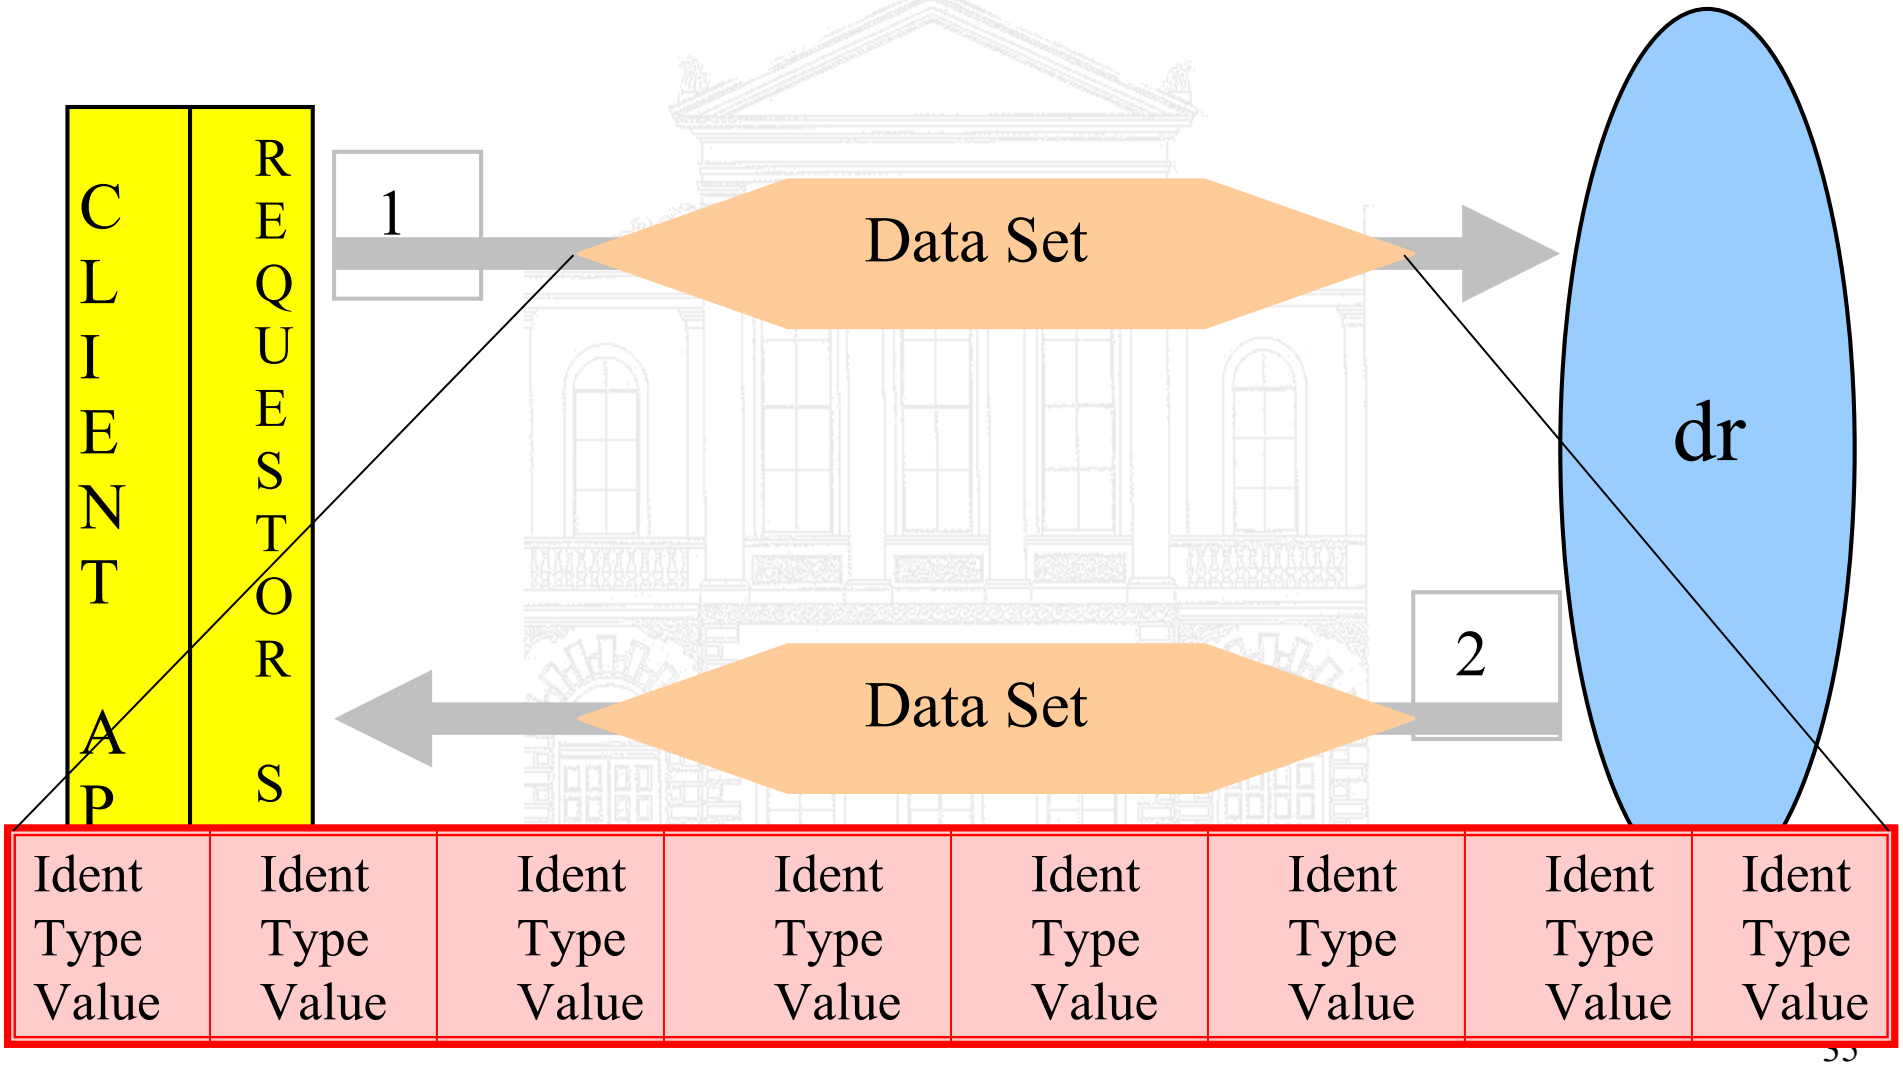

Architecture of Service Interaction

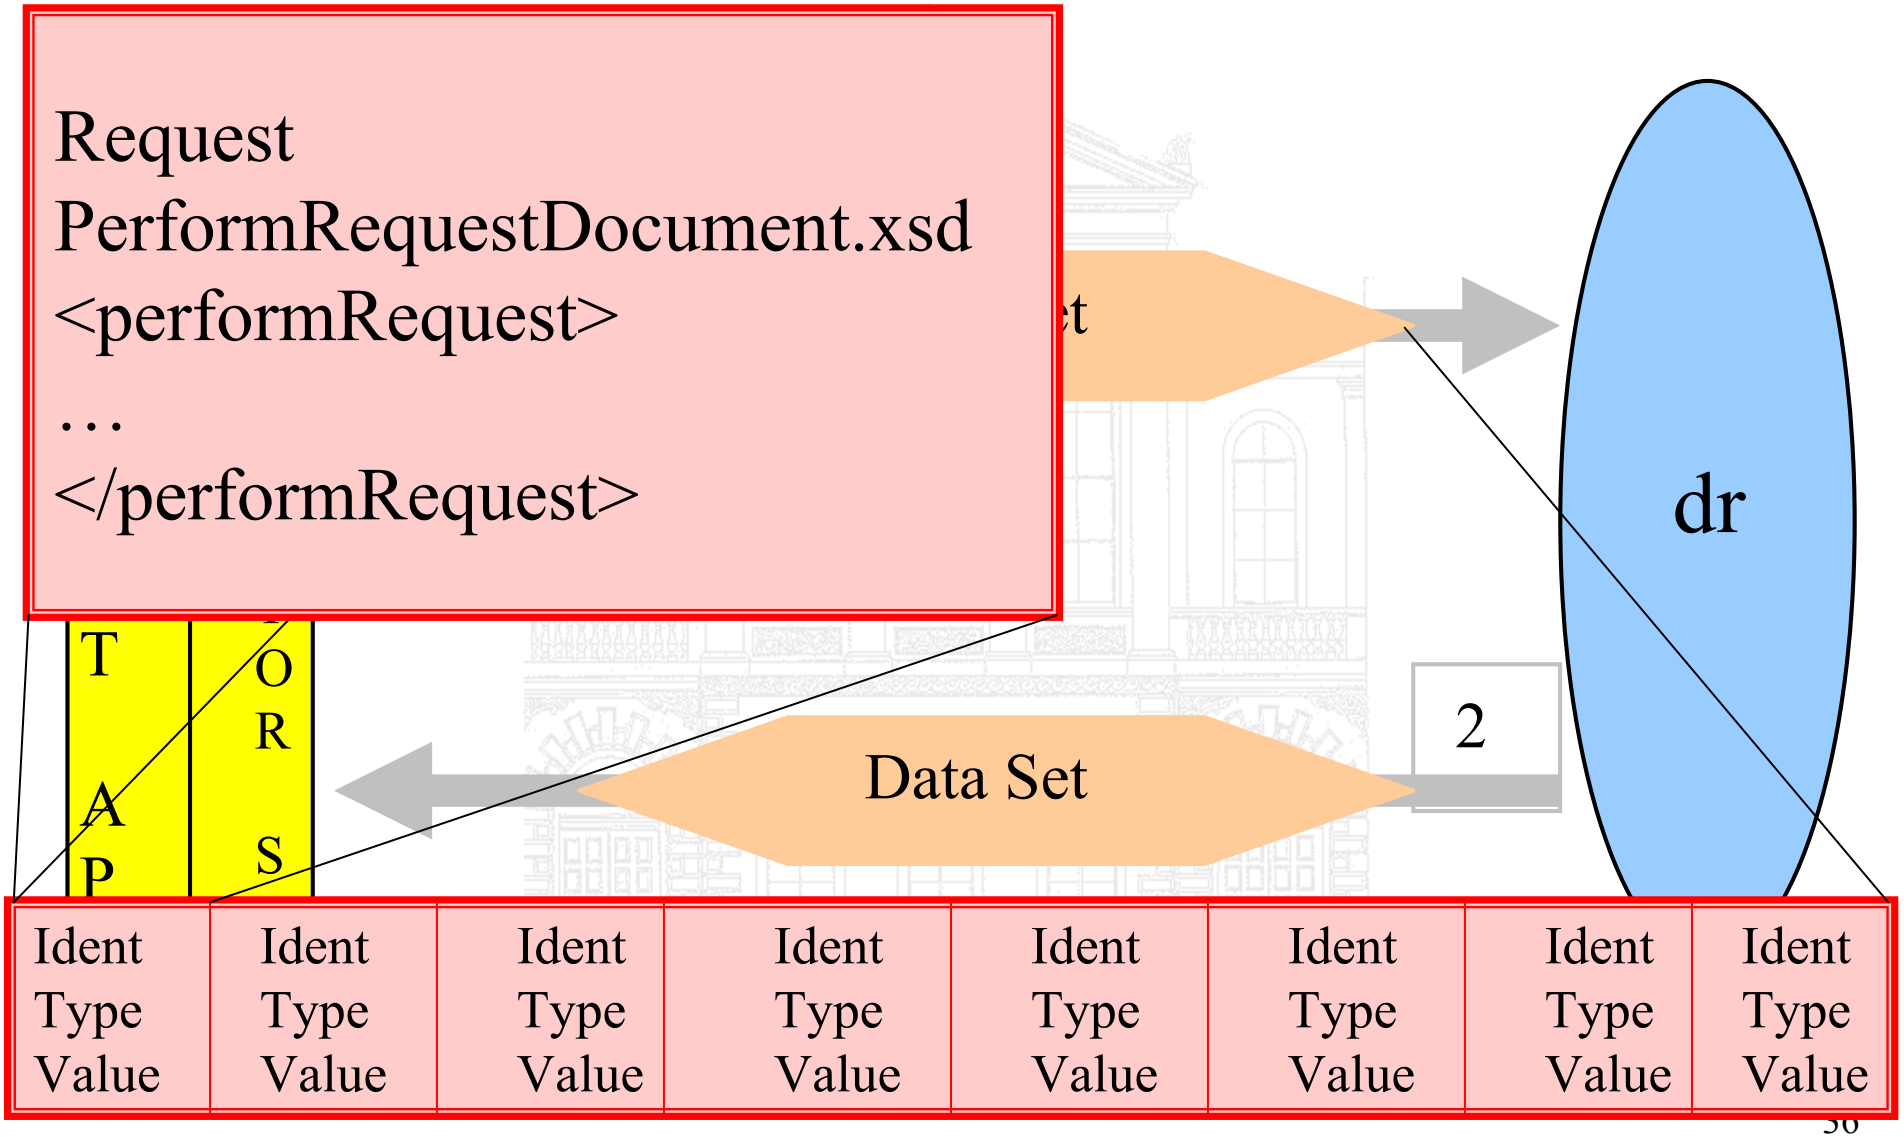

Architecture of Service Interaction

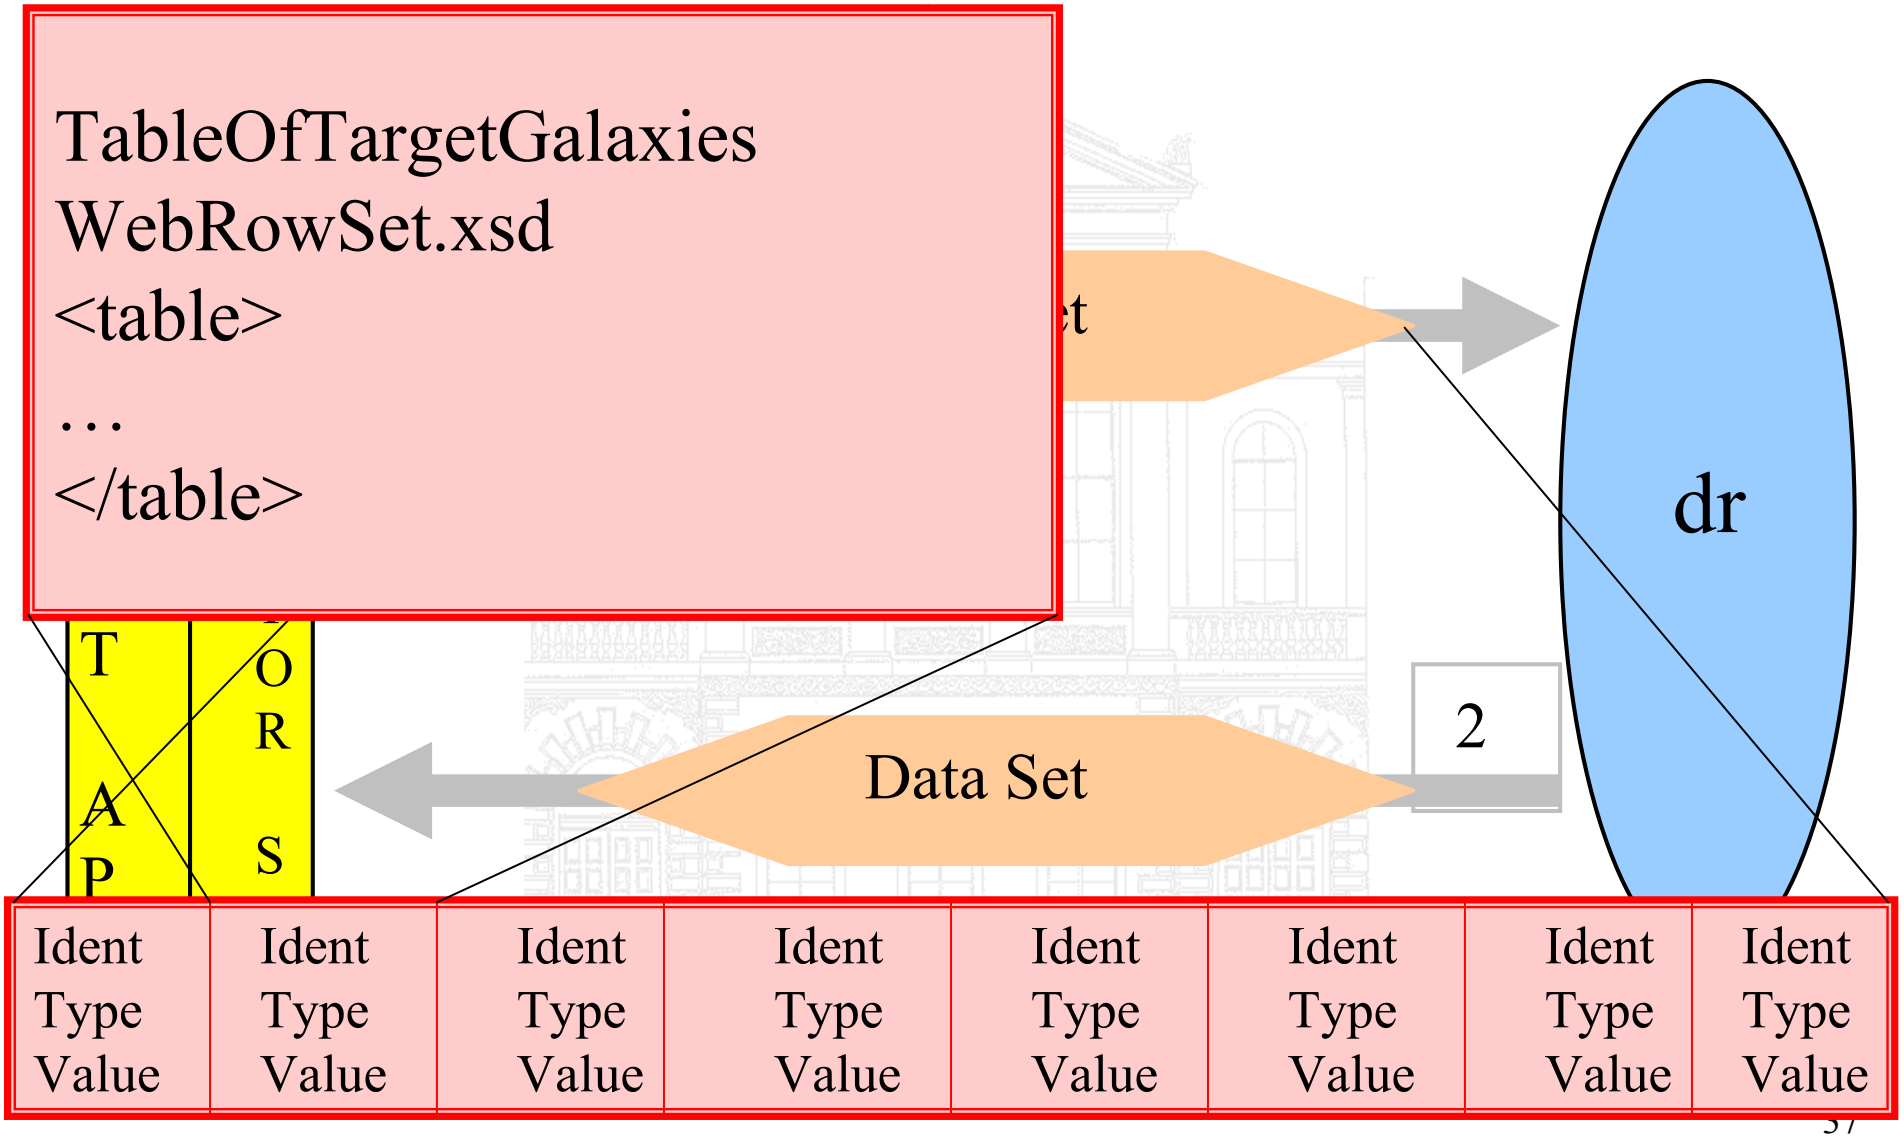

Architecture (2)

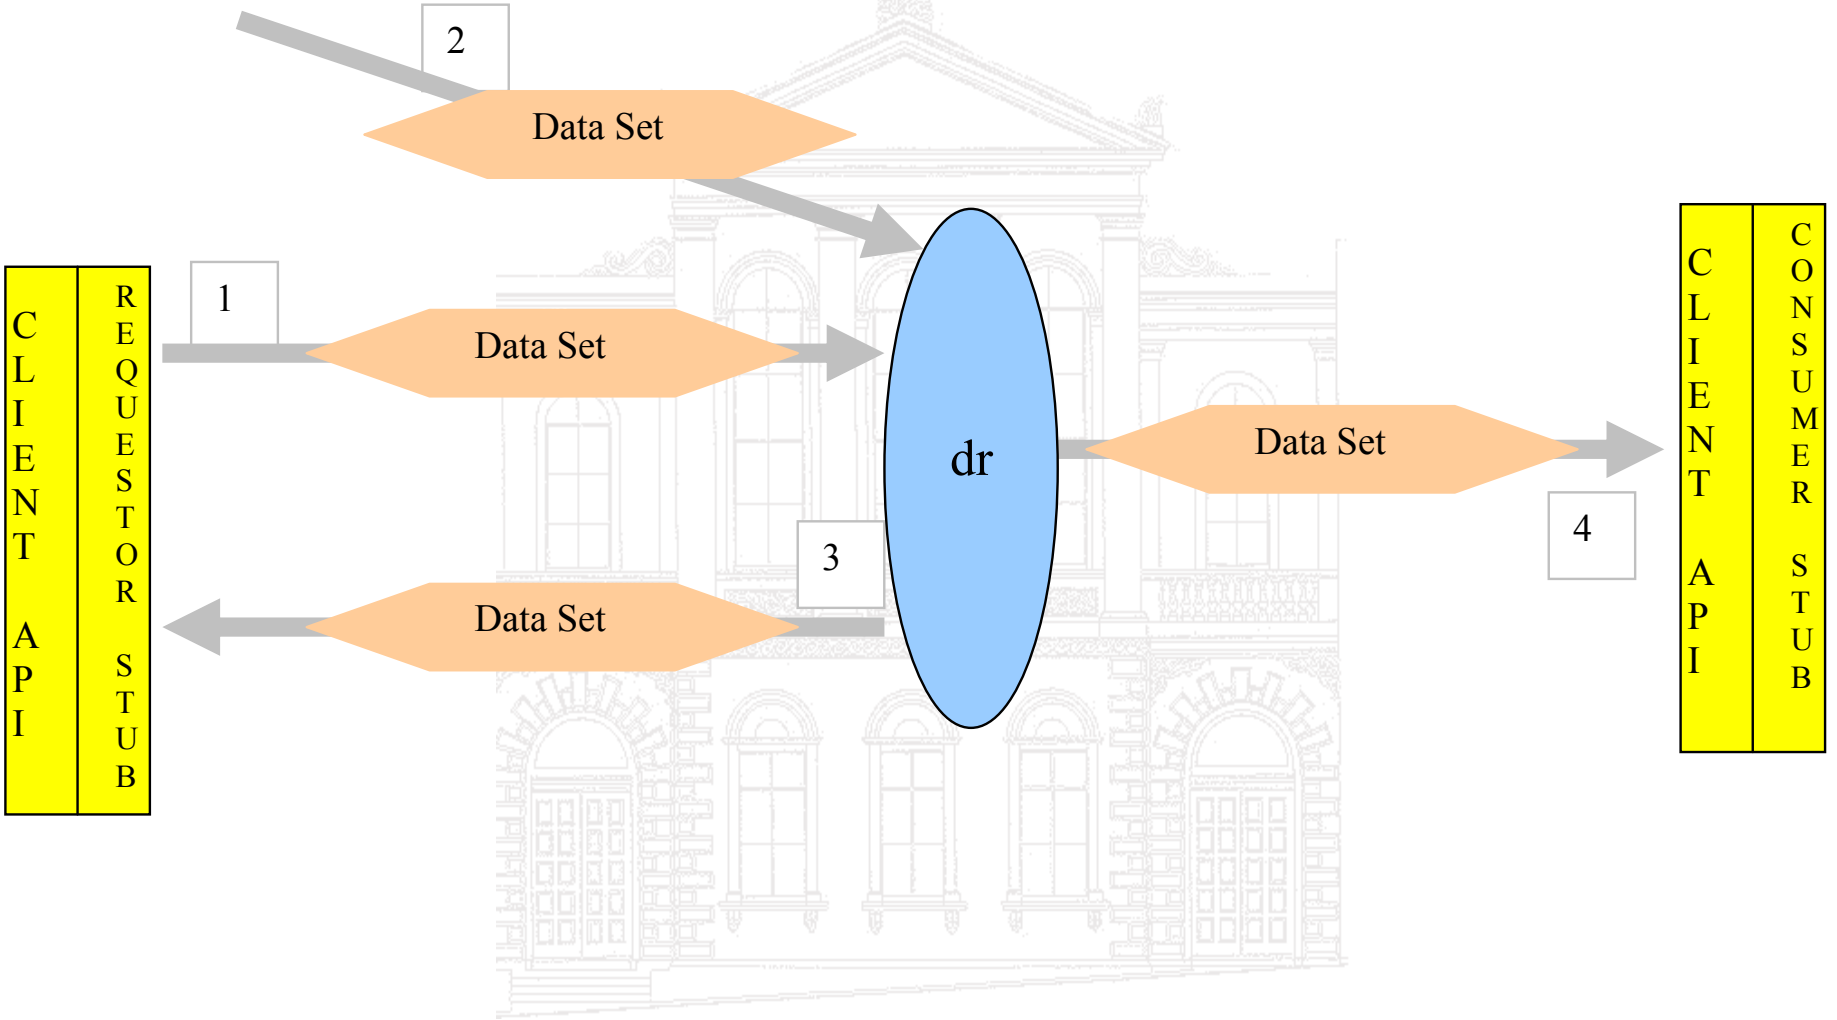

OGSA-DAI Project

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

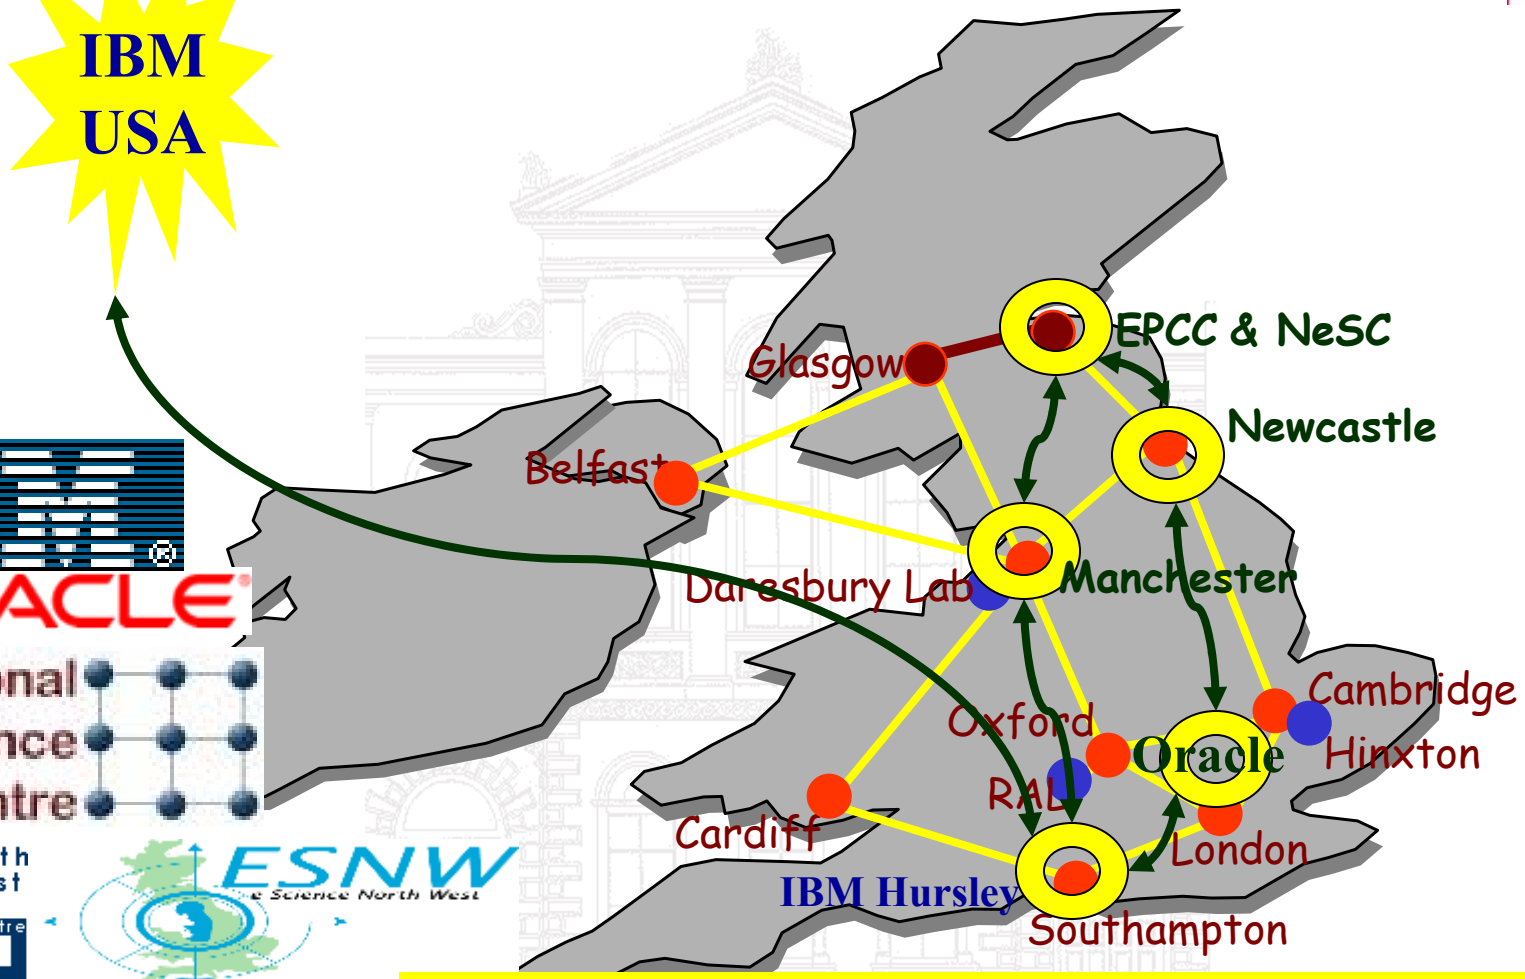

$5 million, 20 months, started February 2002

Additional 24 months, starts October 2003

Infrastructure Architecture

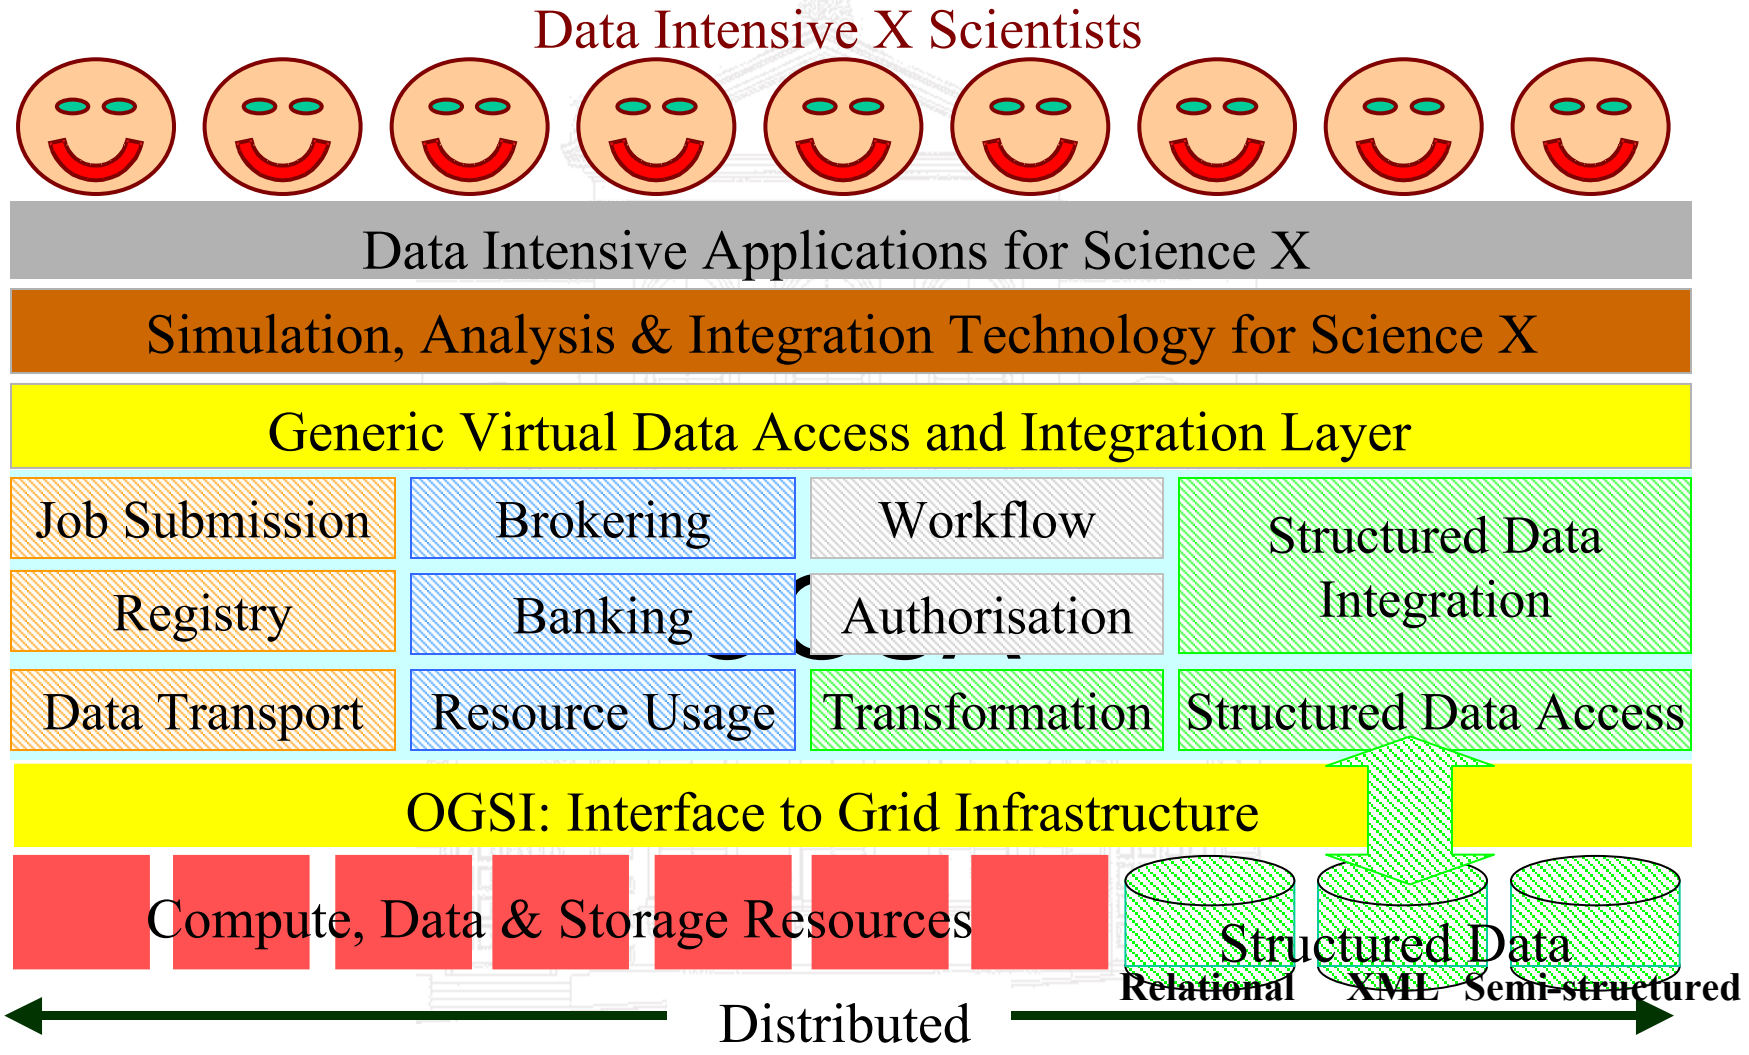

Virtual Integration Architecture

Data Access & Integration Services

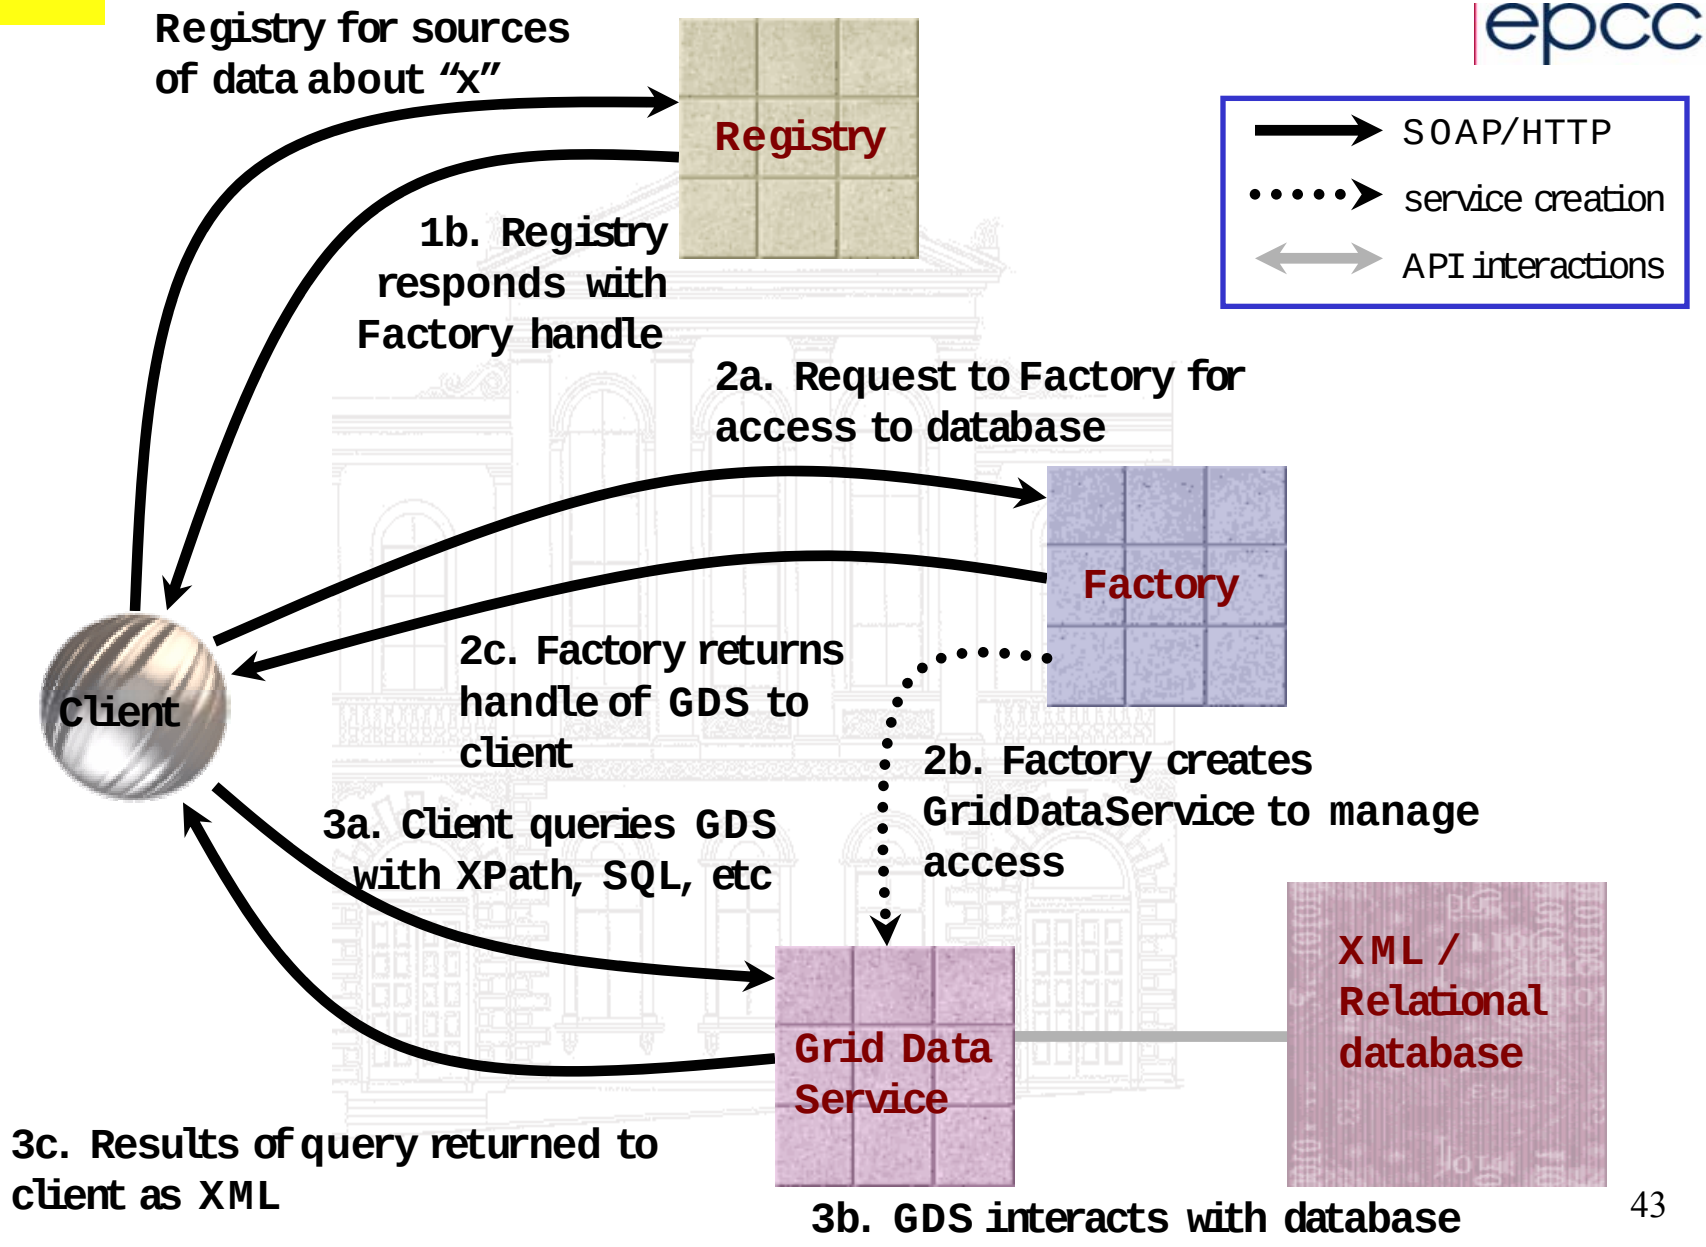

Peering into the Future

Future DAI Services

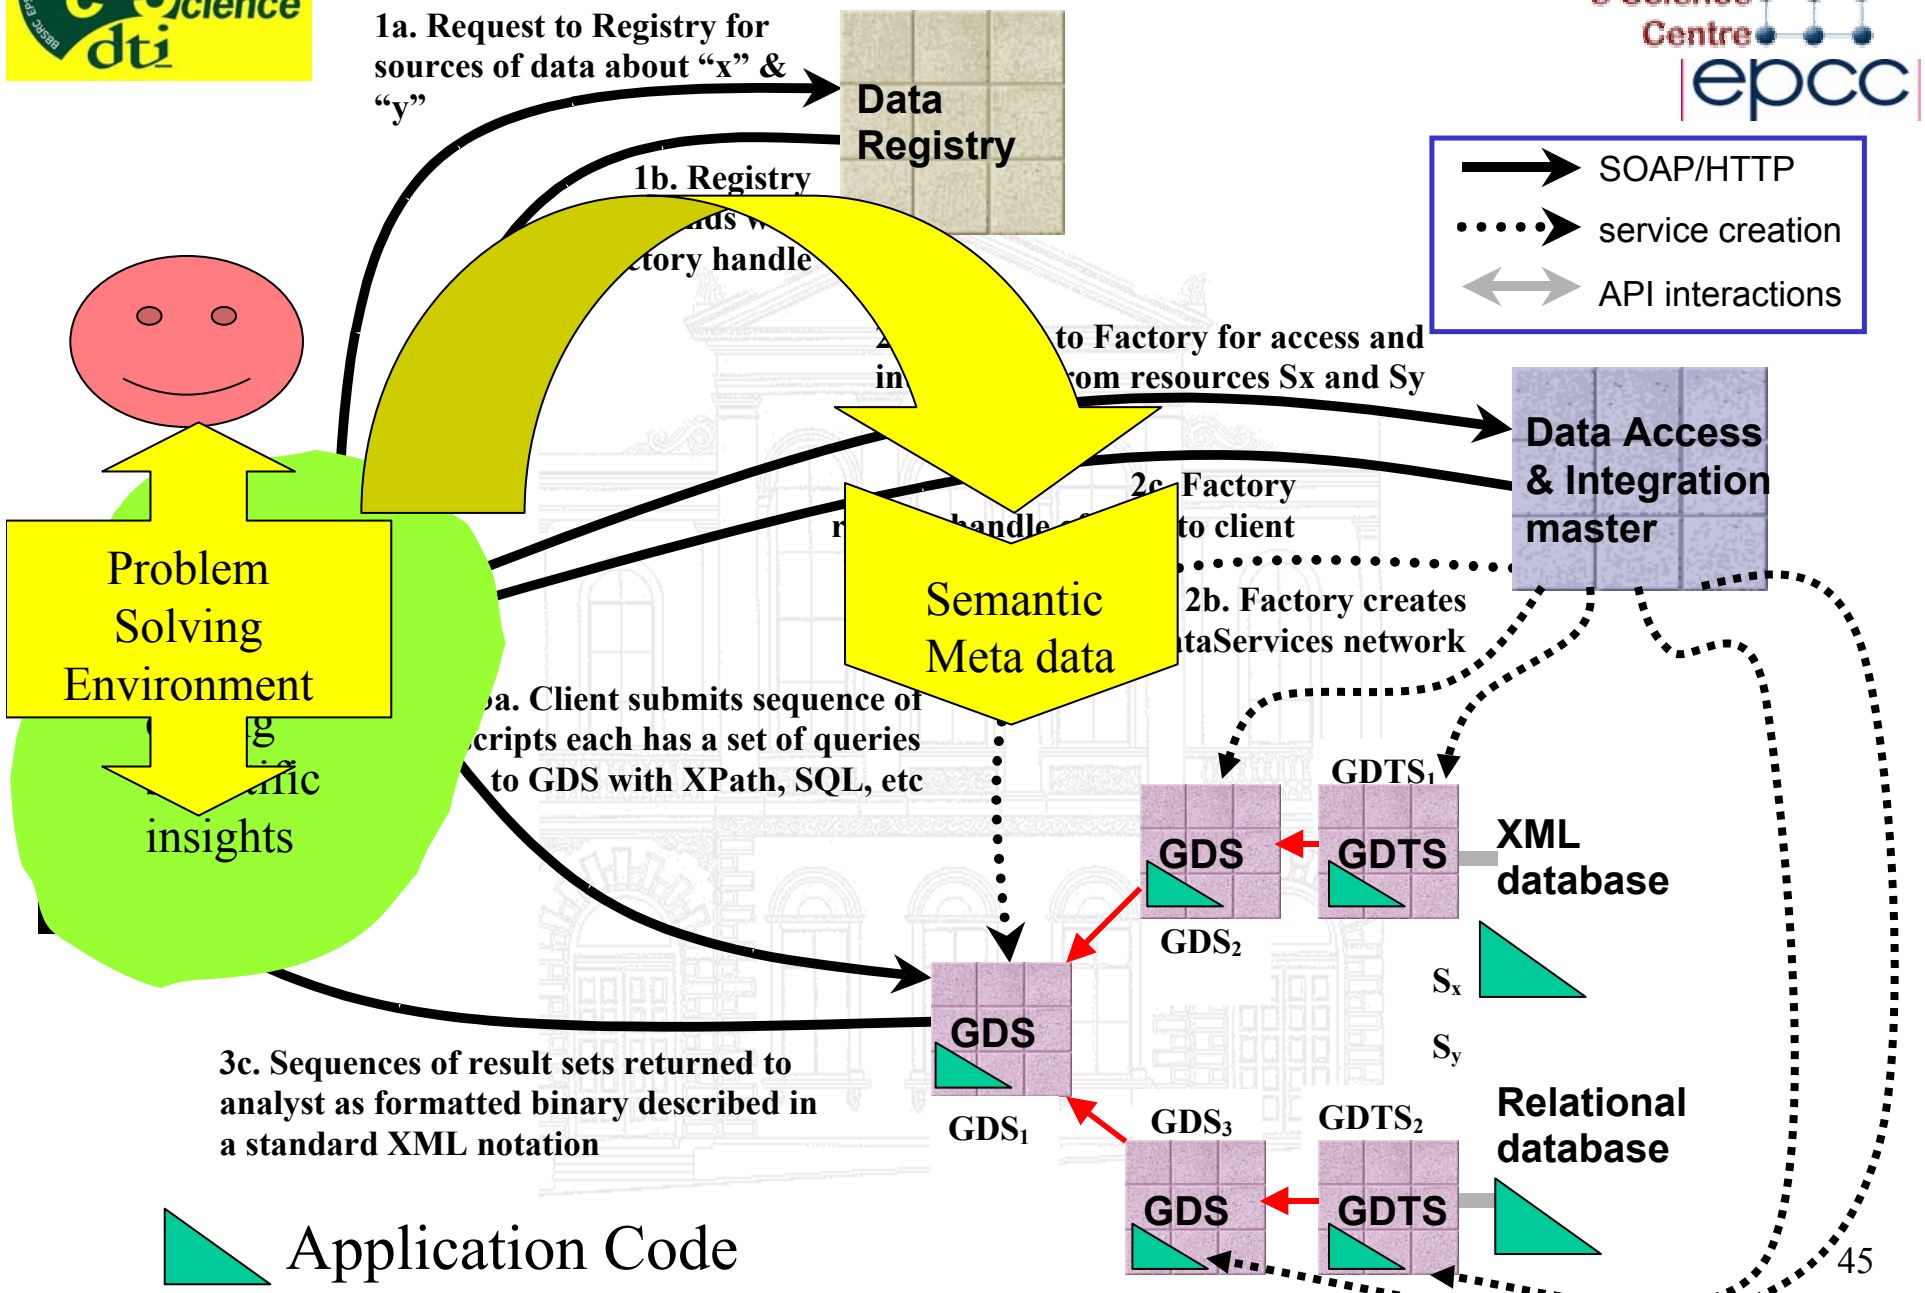

A New World

- What Architecture will Enable Data & Computation Integration?
  - 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?

Take Home Message

- **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