the globus project™
www.globus.org

Open Grid Services as an Enabler of Future Networked Applications

Ian Foster

Argonne National Laboratory
University of Chicago

http://www.mcs.anl.gov/~foster

TERENA Networking Conference, May 20, 2003

Abstract

Grid technologies and infrastructure are designed to support the integration of services and resources within and among enterprises, and thus to allow active collaborations across distributed, multi-organizational collaborations. Recent progress on the Open Grid Services Architecture (OGSA), which integrates Grid technologies with emerging Web services standards, is enabling broad adoption and deployment. I describe the current state and likely evolution of OGSA, and discuss implications for infrastructure and applications.

Partial Acknowledgements

- Open Grid Services Architecture design
  - Carl Kesselman, Karl Czajkowski @ USC/ISI
  - Steve Tuecke @ANL
  - Jeff Nick, Steve Graham, Jeff Frey @ IBM
- Grid services collaborators at ANL
  - Kate Keahey, Gregor von Laszewski
  - Thomas Sandholm, Jarek Gawor, John Bresnahan
- Globus Toolkit R&D also involves many fine scientists & engineers at ANL, USC/ISI, and elsewhere (see www.globus.org)
- Strong links with many EU, UK, US Grid projects
- Support from DOE, NASA, NSF, IBM, Microsoft

Overview

- Problem solving in the 21st century
- Open Grid Services Architecture
- Globus Toolkit v3
- Summary

the globus project™
www.globus.org

Problem Solving in the 21st Century

- Teams organized around common goals
  - Communities: "Virtual organizations"
- With diverse membership & capabilities
  - Heterogeneity is a strength not a weakness
- And geographic and political distribution
  - No location/organization possesses all required skills and resources
- Must adapt as a function of the situation
  - Adjust membership, reallocate responsibilities, renegotiate resources

For Example: Global Knowledge Communities

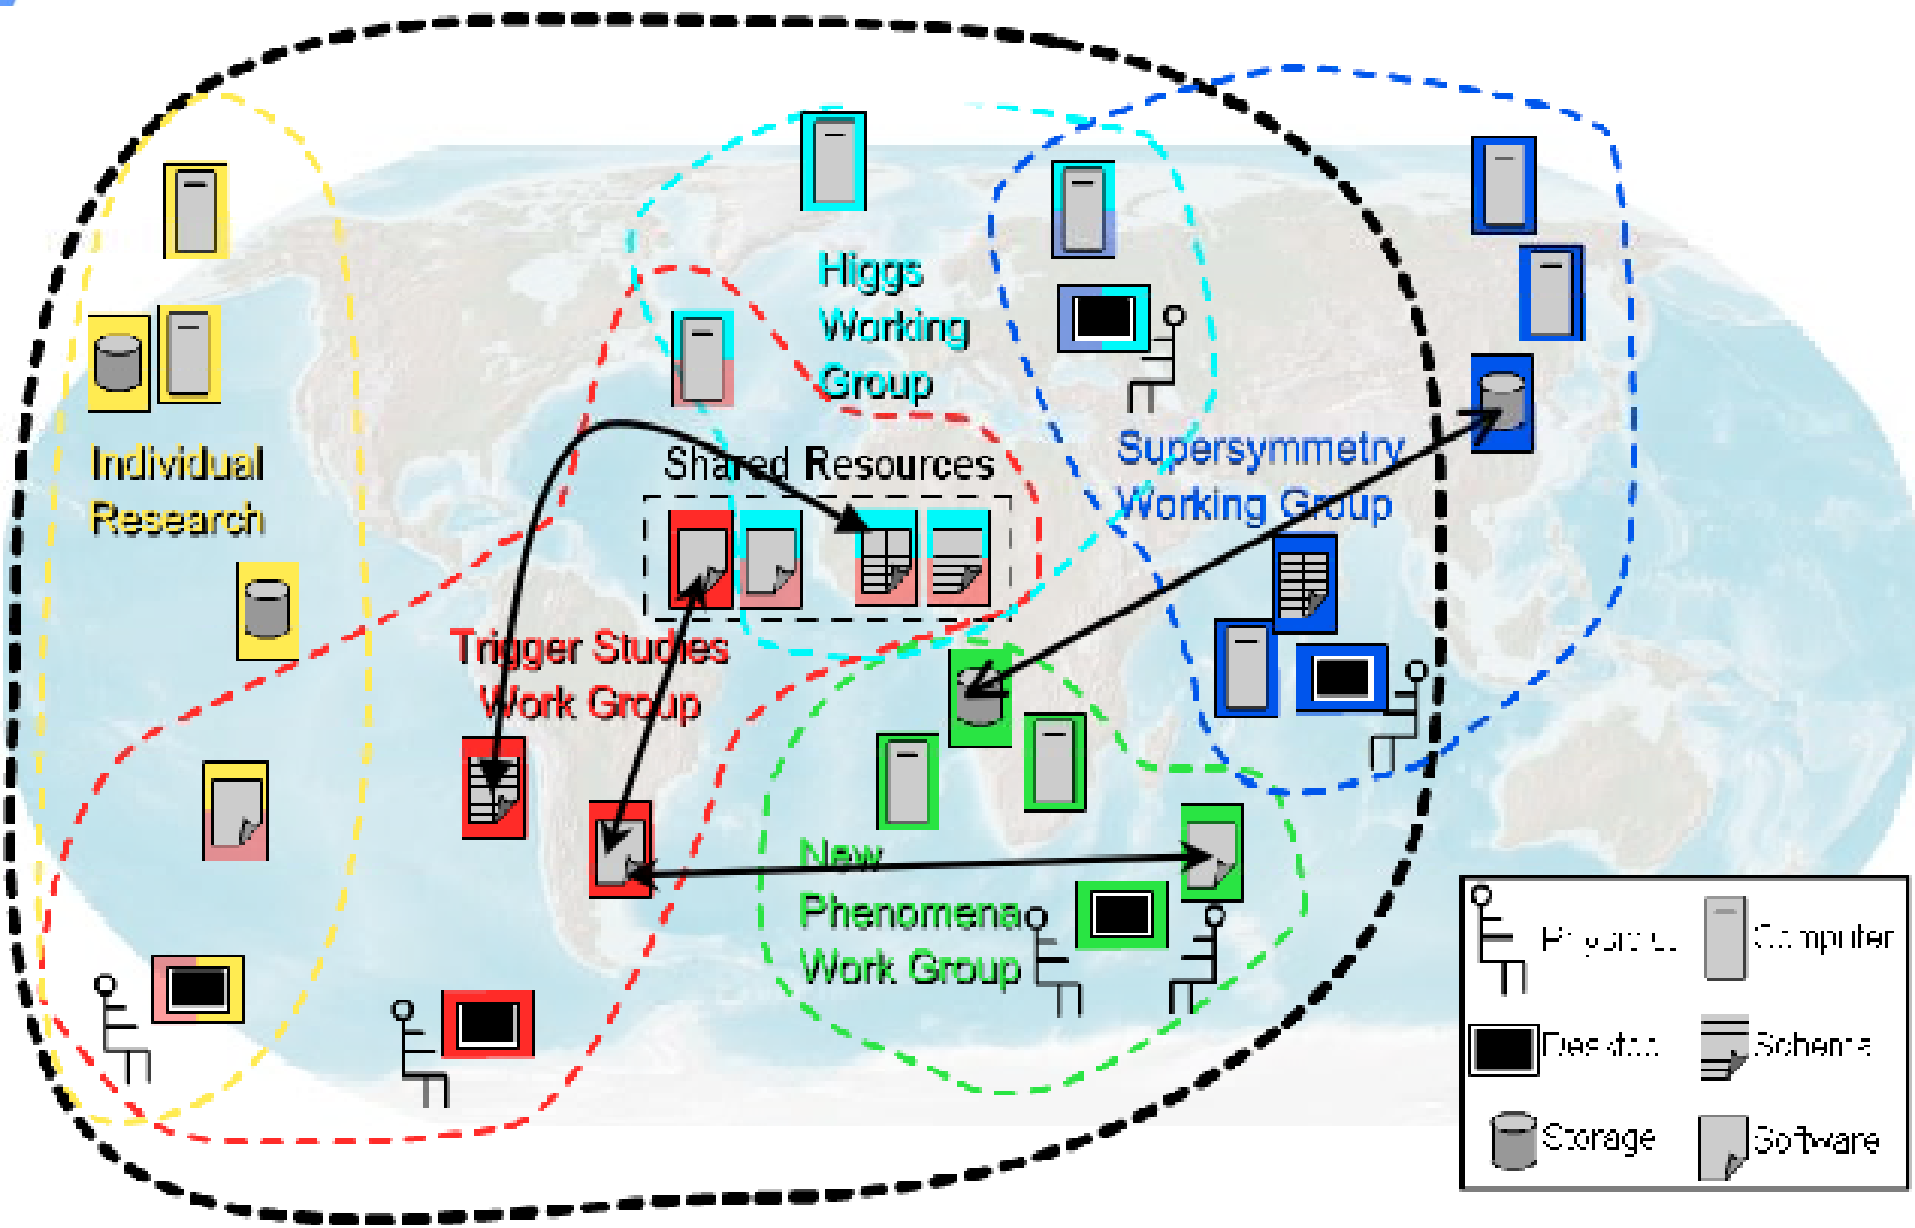

New Opportunities Demand New Technology

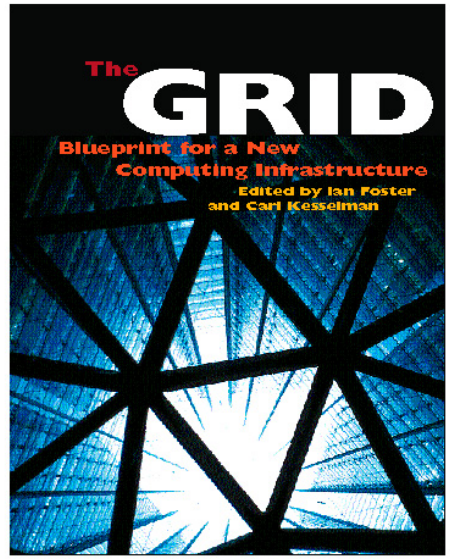

"Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations"

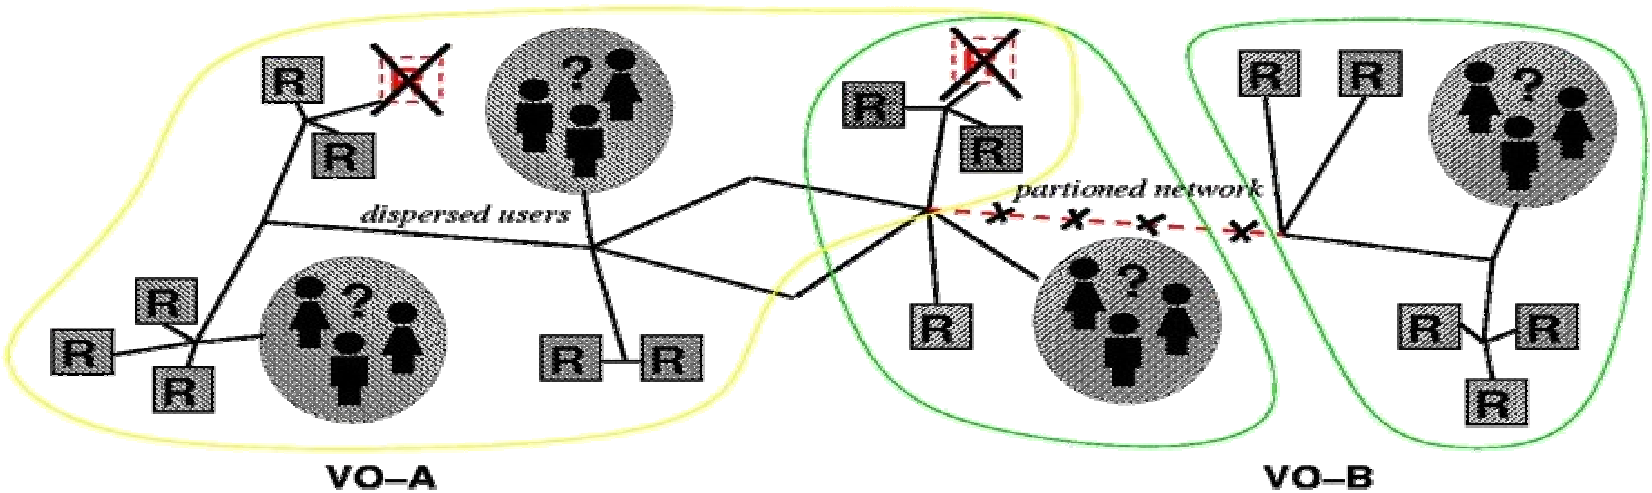

"*When the network is as fast as the computer's internal links, the machine disintegrates across the net into a set of special purpose appliances*" (George Gilder)

the globus project™
www.globus.org

Taking Sharing to the Next Level

- Sharing of communication
  - Telephones, mailing lists, collaboration tools
- Sharing of data and knowledge
  - Web, semantic web
- What about the rest of the infrastructure?
  - Services, computers, programs, sensors, …

Existing Technologies are Helpful, but Not Complete Solutions

- Peer-to-peer technologies
  - Limited scope and mechanisms
- Enterprise-level distributed computing
  - Limited cross-organizational support
- Databases
  - Vertically integrated solutions
- Web services
  - Not dynamic
- Semantic web
  - Limited focus

What's Missing is Support for …

- Sharing & integration of resources, via
  - Discovery
  - Provisioning
  - Access (computation, data, …)
  - Security
  - Policy
  - Fault tolerance
  - Management
- In dynamic, scalable, multi-organizational settings

Enter the Grid

- Infrastructure ("middleware") for establishing, managing, and evolving multi-organizational federations
  - Dynamic, autonomous, domain independent
  - On-demand, ubiquitous access to computing, data, and services
- Mechanisms for creating and managing workflow within such federations
  - New capabilities constructed dynamically and transparently from distributed services
  - Service-oriented, virtualization

the globus project™
www.globus.org

The Emergence of Open Grid Standards

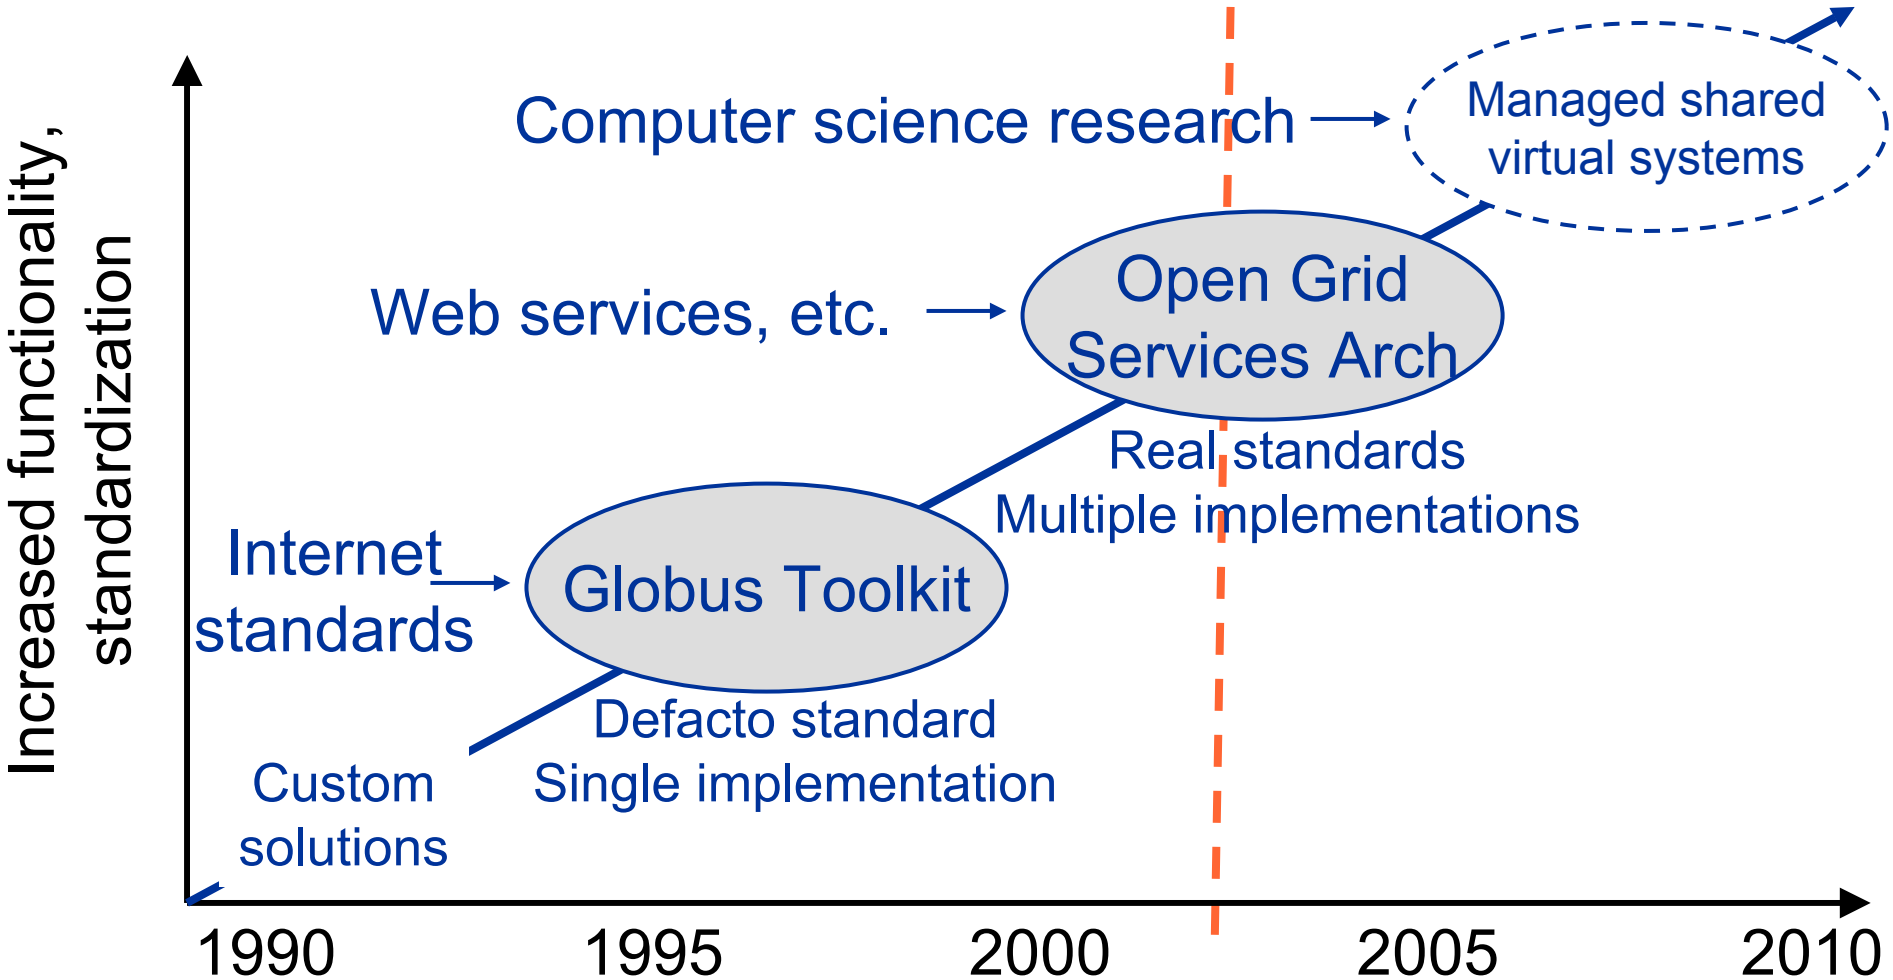

Overview

- Problem solving in the 21st century
- **Open Grid Services Architecture**
- Globus Toolkit v3
- Summary

the globus project™
www.globus.org

Open Grid Services Architecture

- Service-oriented architecture
  - Key to virtualization, discovery, composition, local-remote transparency
- Leverage industry standards
  - Internet, Web services
- Distributed service management
  - A "component model for Web services"
- A framework for the definition of composable, interoperable services

"The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration", Foster, Kesselman, Nick, Tuecke, 2002

the globus project™
www.globus.org

Web Services

- XML-based distributed computing technology
- Web service = a server process that exposes typed ports to the network
- Described by the Web Services Description Language, an XML document that contains
  - Type of message(s) the service understands & types of responses & exceptions it returns
  - "Methods" bound together as "port types"
  - Port types bound to protocols as "ports"
- A WSDL document completely defines a service and how to access it

OGSA Structure

- A standard substrate: the Grid service
  - Standard interfaces and behaviors that address key distributed system issues
  - A refactoring and extension of the Globus Toolkit protocol suite
- … supports standard service specifications
  - Resource management, databases, workflow, security, diagnostics, etc., etc.
  - Target of current & planned GGF efforts
- … and arbitrary application-specific services based on these & other definitions

the globus project™
www.globus.org

Open Grid Services Infrastructure

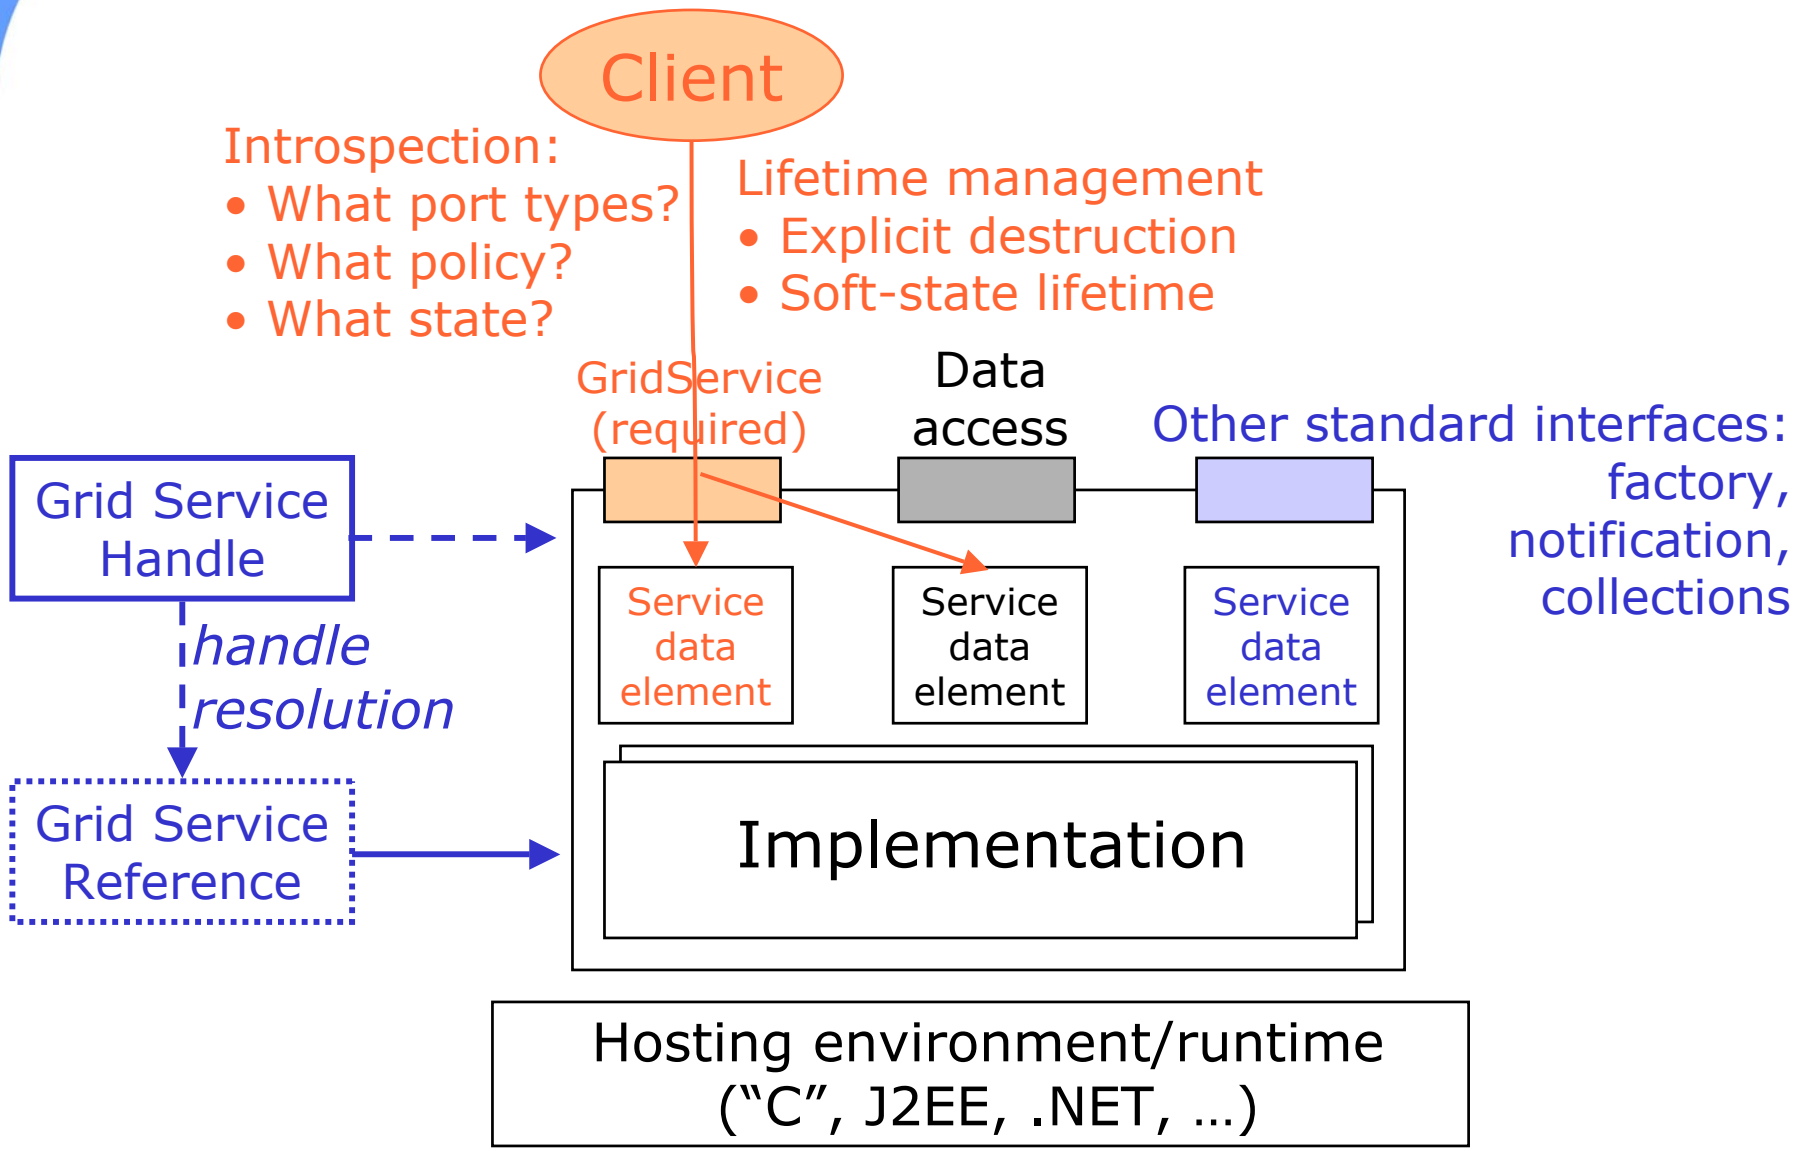

Open Grid Services Infrastructure

GWD-R (draft-ggf-ogsi- gridservice-23)
Open Grid Services Infrastructure (OGSI)
http://www.ggf.org/ogsi-wg

Editors:
S. Tuecke, ANL
K. Czajkowski, USC/ISI
I. Foster, ANL
J. Frey, IBM
S. Graham, IBM
C. Kesselman, USC/ISI
D. Snelling, Fujitsu Labs
P. Vanderbilt, NASA
February 17, 2003

Open Grid Services Infrastructure (OGSI)

the globus project™
www.globus.org

Example: Reliable File Transfer Service

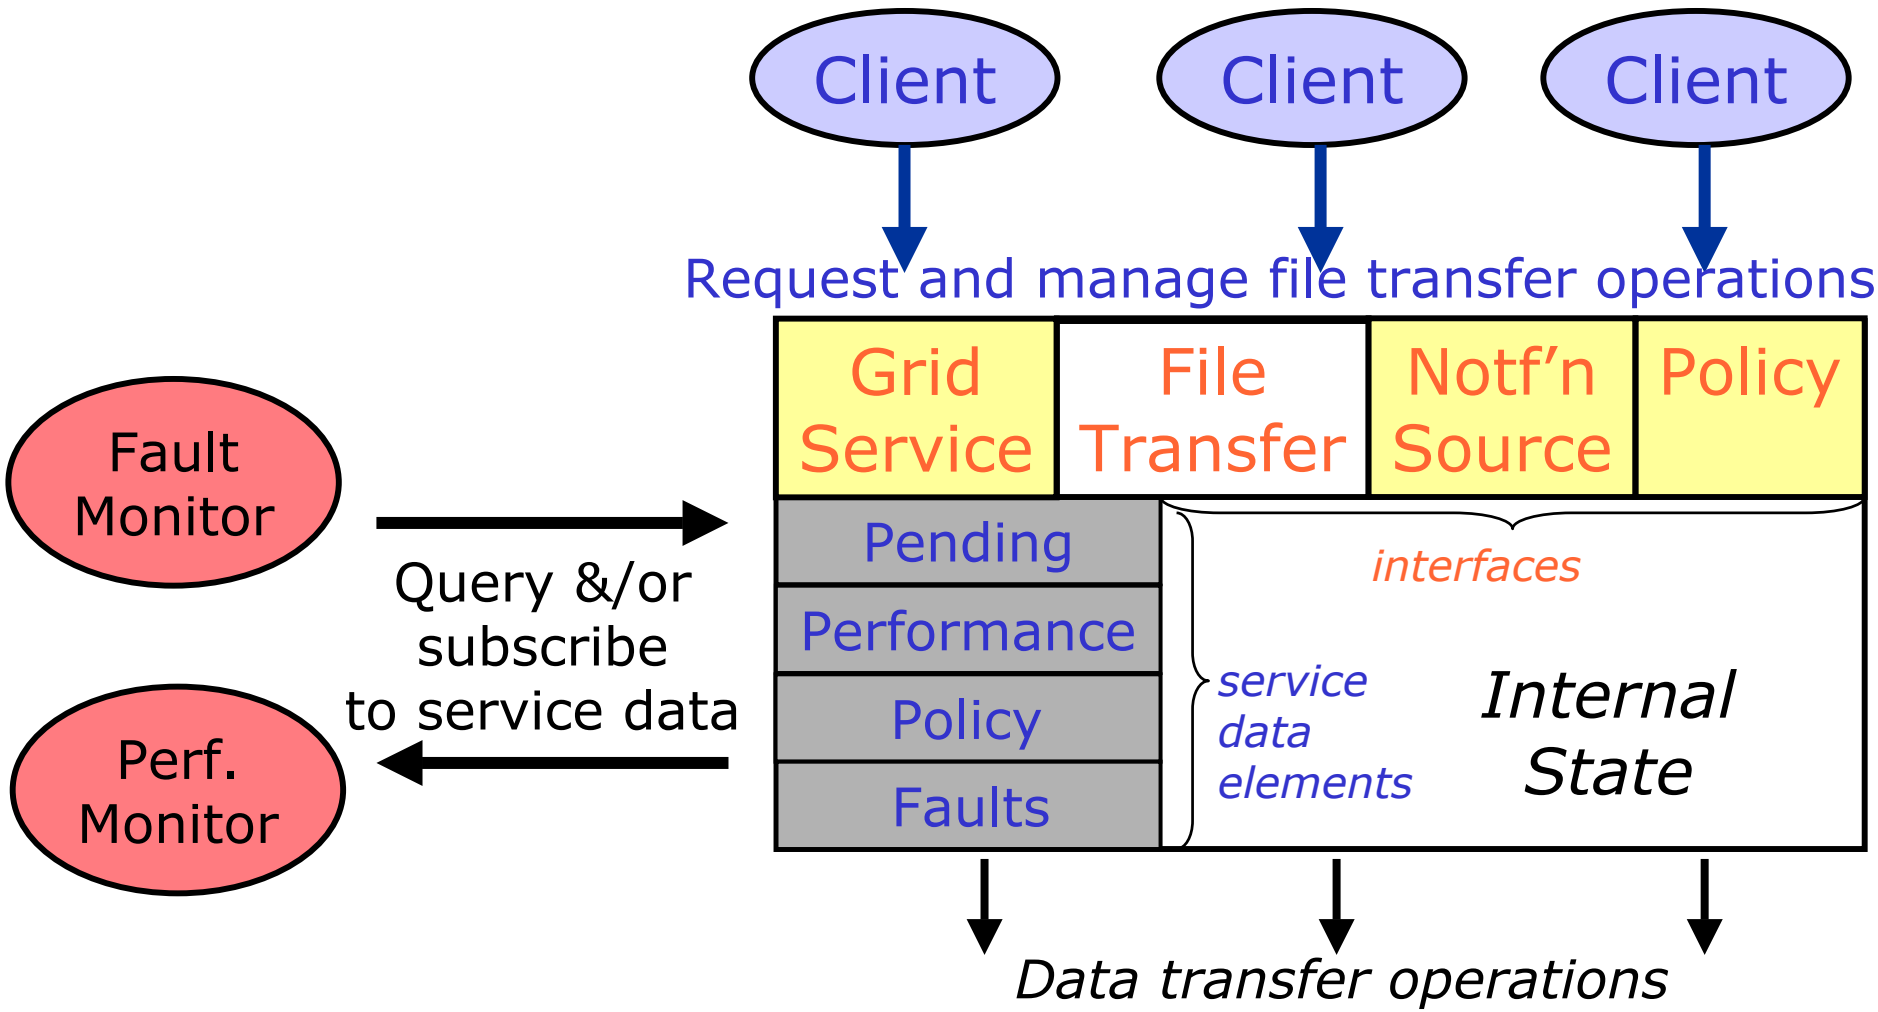

the globus project™
www.globus.org

Open Grid Services Architecture

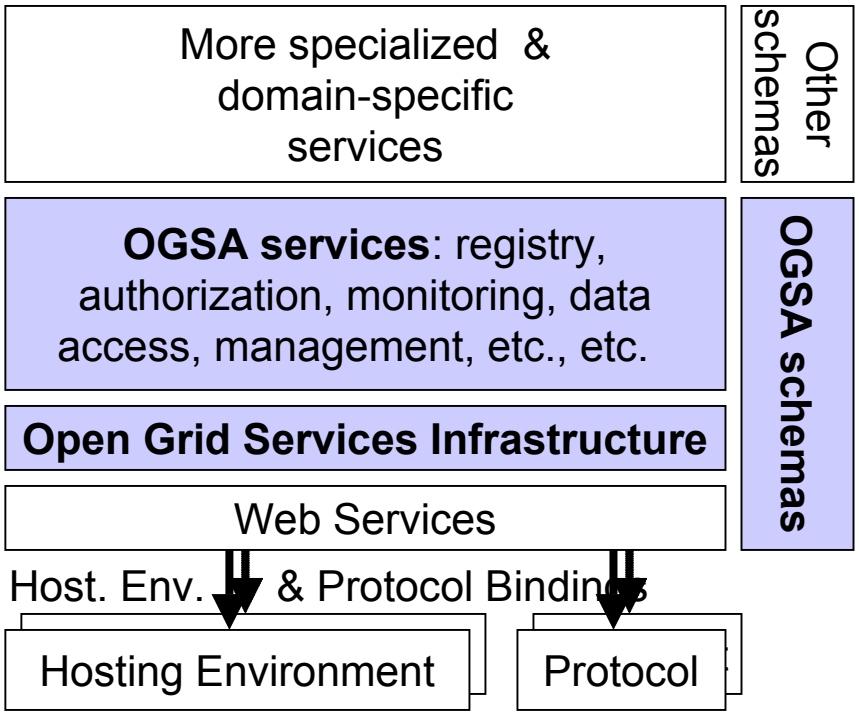

- Data access and integration
- Security
- SLA negotiation
- Manageability
- ...

GWD-R (draft-ggf-ogsa-platform-3)
Open Grid Services Architecture Platform
http://www.ggf.org/ogsa-wg

Editors:
I. Foster, Argonne & U.Chicago
D. Gannon, Indiana U.

Grid Security Services

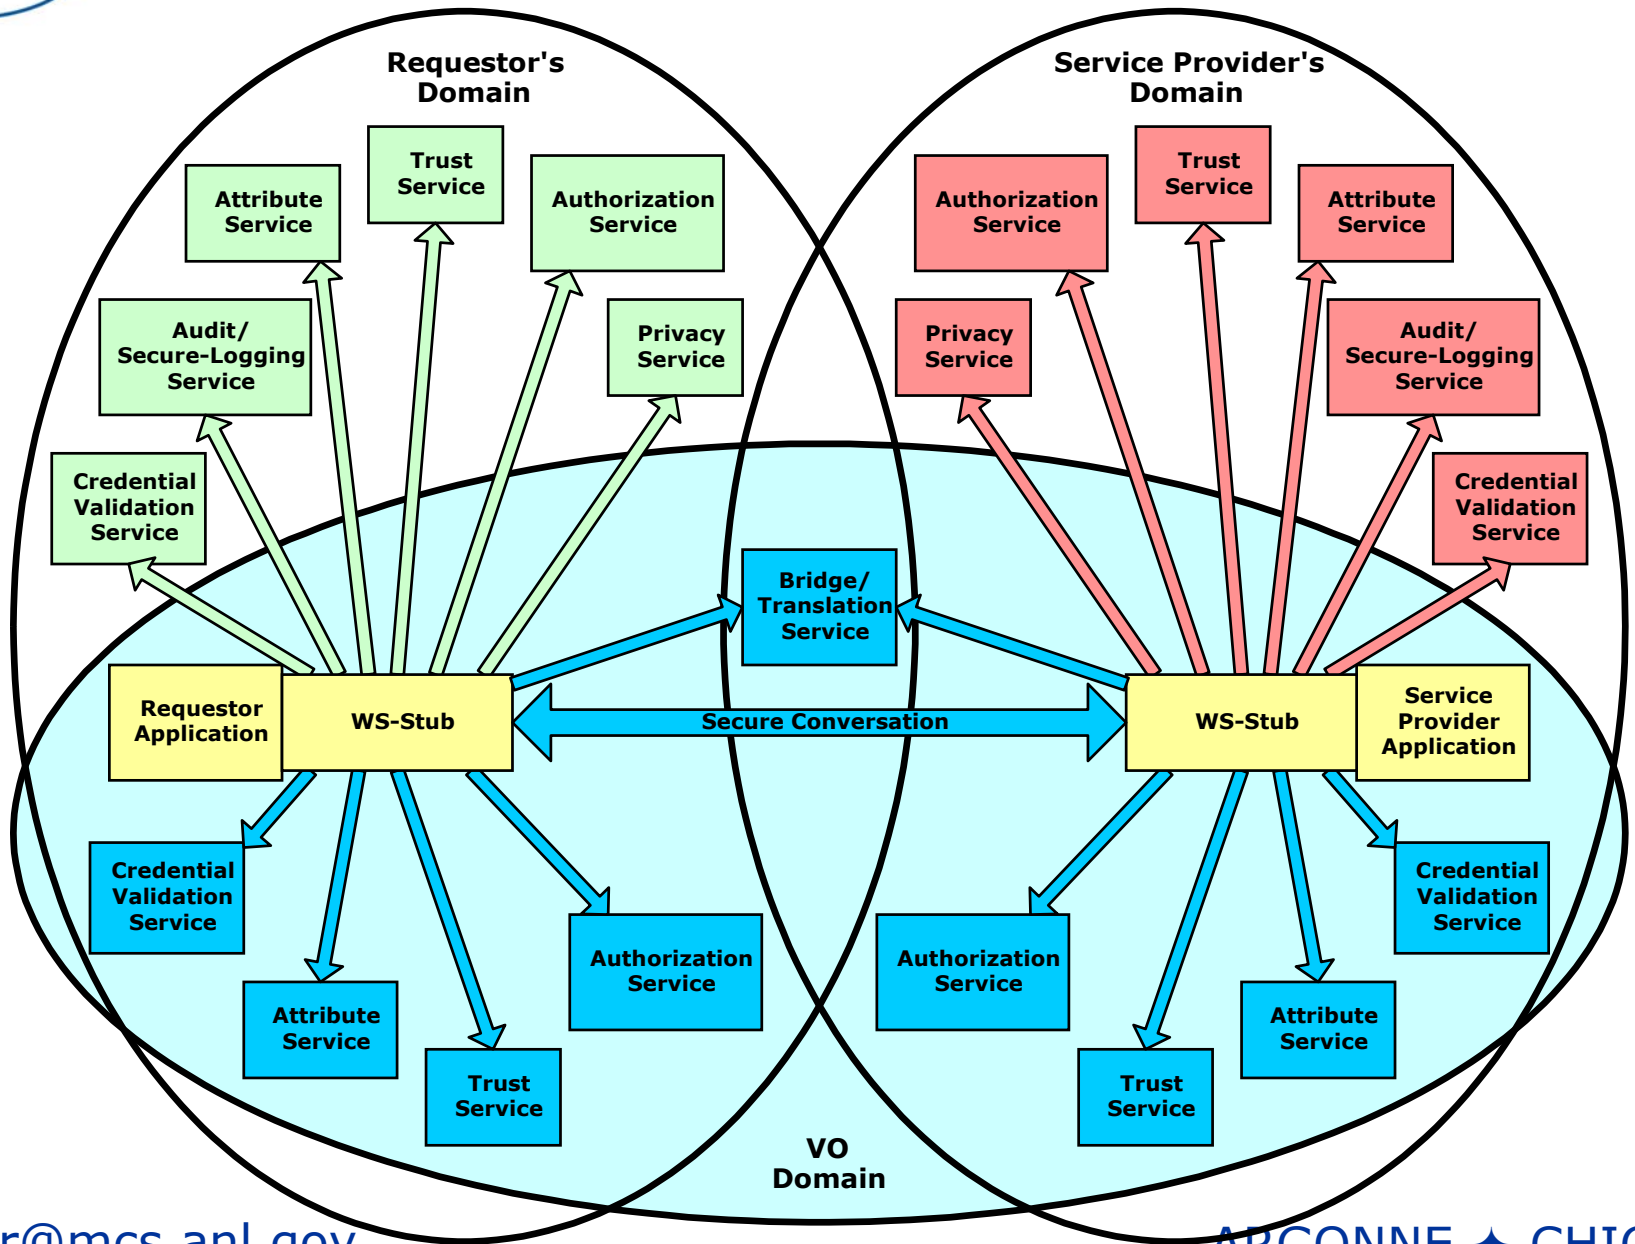

OGSA Manageability

- Management is a high-priority OGSA effort
  - GGF Common Management Model (CMM) WG
  - Ellen Stokes (IBM) co-chair
- Goal:
  - Define standard schema and interfaces for a manageable resource (modeled as a service)
- And also:
  - Allow existing models to be used & exposed
  - Leverage CIM schema when applicable
  - Define how CMM interfaces can integrate with higher-level interfaces (e.g., provisioning)

Common Management Model

- A manageable resource is a Grid service, thus
  - Global resource names: Grid service handles
  - State data modeling + access: SDEs
  - Lifetime management
  - Service Group for grouping resources
  - Interface definition language: WSDL
- Plus additional schema & operations
  - Standard manageable resource SDE schema
  - Interfaces for extensible lifecycle and relationship management

⇒ BaseManageableResource interface

Base Manageable Port Types

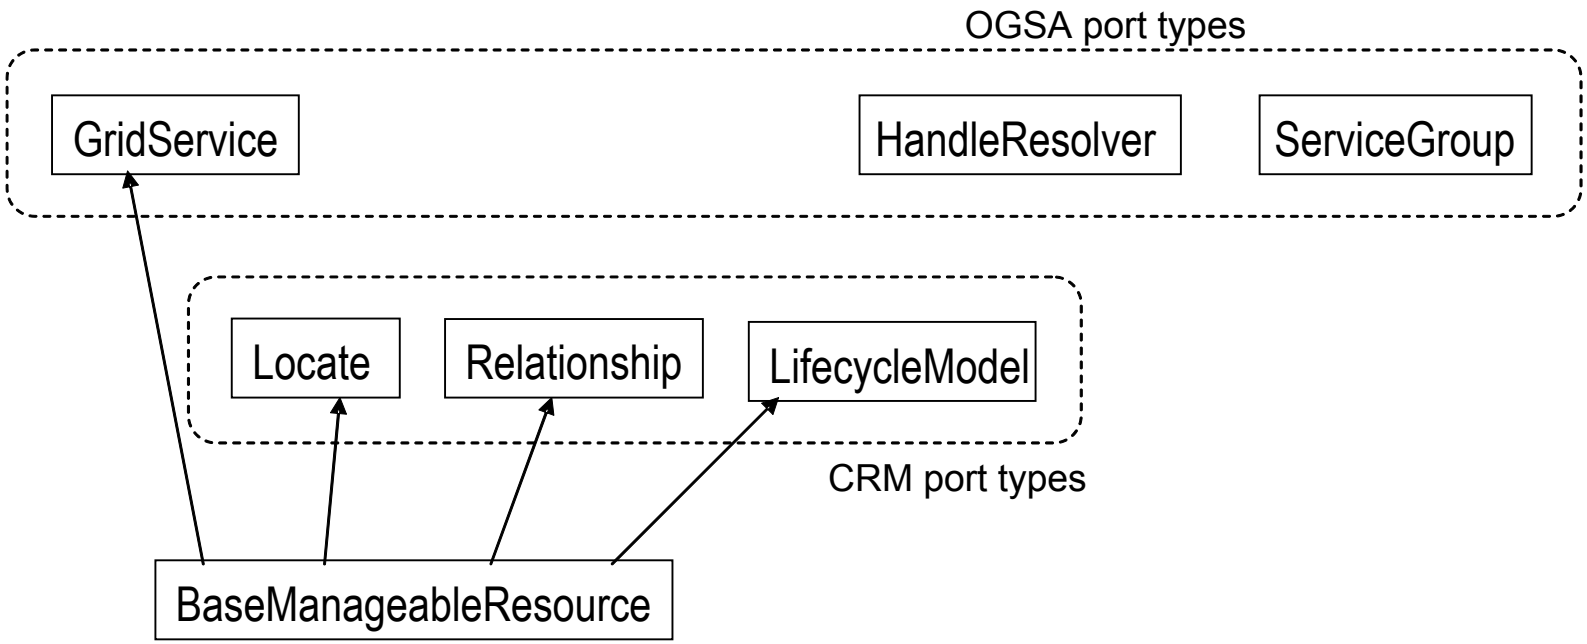

CMM Schema

- WSDL (open content model) & XSD describe resource's manageable attributes (as SDEs)
- Models are CIM-based where applicable
  - Re-factored for service efficiency; appropriate for higher-level management applications
  - Additional XML attributes: change control, measuring, lifecycle
  - Additional XML data types
- CIM as basis for schemas: but may modify
  - E.g. use constructs from XML/XSD where similar ones exist in CIM

the globus project™
www.globus.org

"Use CIM Models Where Applicable"

- Class is port type, properties of class are port type service data, methods of class are port type operations
  - Some refactoring of classes -> portTypes
- Express in WSDL/GSDL as grid service
  - Managed resource port type from which other resource port types are derived
  - Mix in the base GridService port type
  - Mix in other CMM port types as needed (Identity, relationship, lifecycleState)

OGSA Misconceptions

- OGSA means you have to code in Java
  - No: C client bindings now, C server side eventually (but not needed for current apps)
- OGSA means all programs must be services
  - No: You can write services if you want, but other behaviors are supported: e.g., GT3 supports GT2 GRAM, GridFTP, …, …
- OGSA is a silver bullet for distributed and collaborative computing
  - No, it makes some things easier, but it's only interfaces and behaviors, after all!

Overview

- Problem solving in the 21st century
- Open Grid Services Architecture
- **Globus Toolkit v3**
- Summary

Globus Toolkit v3 (GT3) Open Source OGSA Technology

- Implements OGSI interfaces
- Supports primary GT2 interfaces
  - GRAM, GridFTP, GSI
  - High degree of backward compatibility
- Multiple platforms & hosting environments
  - J2EE, Java, C, .NET, Python
- New services
  - SLA negotiation, service registry, community authorization, data management, …
- Broad & growing adoption and contributions

the globus project™
www.globus.org

Globus Toolkit Contributors: GT2

- Grid Packaging Technology (GPT) — NCSA
- Persistent GRAM Jobmanager — Condor
- GSI/Kerberos interchangeability — Sandia
- Documentation — NASA, NCSA
- Ports — IBM, HP, Sun, SDSC, …
- MDS stress testing — EU DataGrid
- Support — IBM, Platform, UK eScience
- Testing and patches — Many!
- Interoperable tools — Many!
- $$ — DARPA, DOE, NSF, NASA, Microsoft, EU

the globus project™
www.globus.org

Globus Toolkit Contributors: GT3

- Replica location service — EU DataGrid
- Python hosting environment — LBNL
- Data access & integration — UK eScience
- Data mediation services — SDSC
- Tooling, Xindice, JMS — IBM
- …
- …
- …

GT2 Evolution To GT3

- What happened to the GT2 key protocols?
  - Security: Adapting X.509 proxy certs to integrate with emerging WS standards
  - GRIP/LDAP/MDS: Abstractions integrated into OGSI as serviceData
  - GRAM: ManagedJobFactory and related service definitions
  - GridFTP: Unchanged in 3.0, but will evolve into OGSI-compliant service in 2003
- Also rendering collective services in terms of OGSI: RFT, RLS, CAS, etc.

the globus project™
www.globus.org

Example: Reliable File Transfer Service

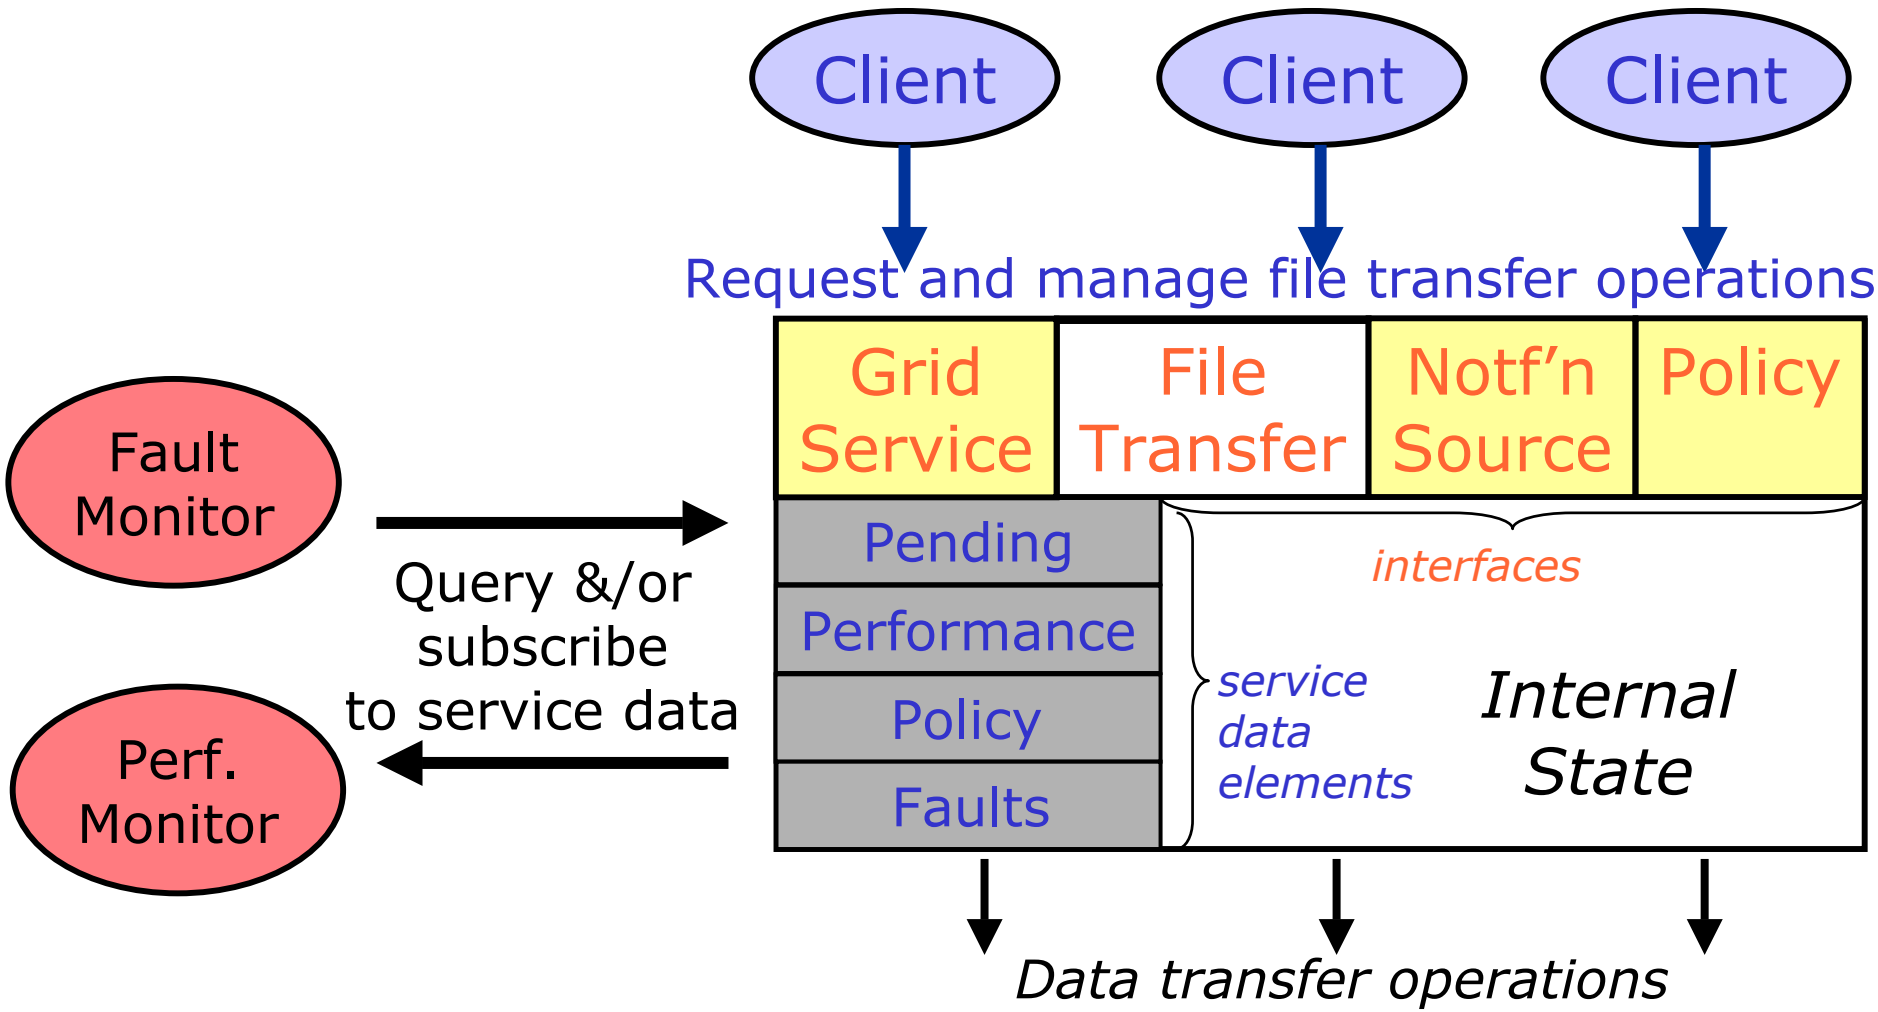

The Grid Technology Repository

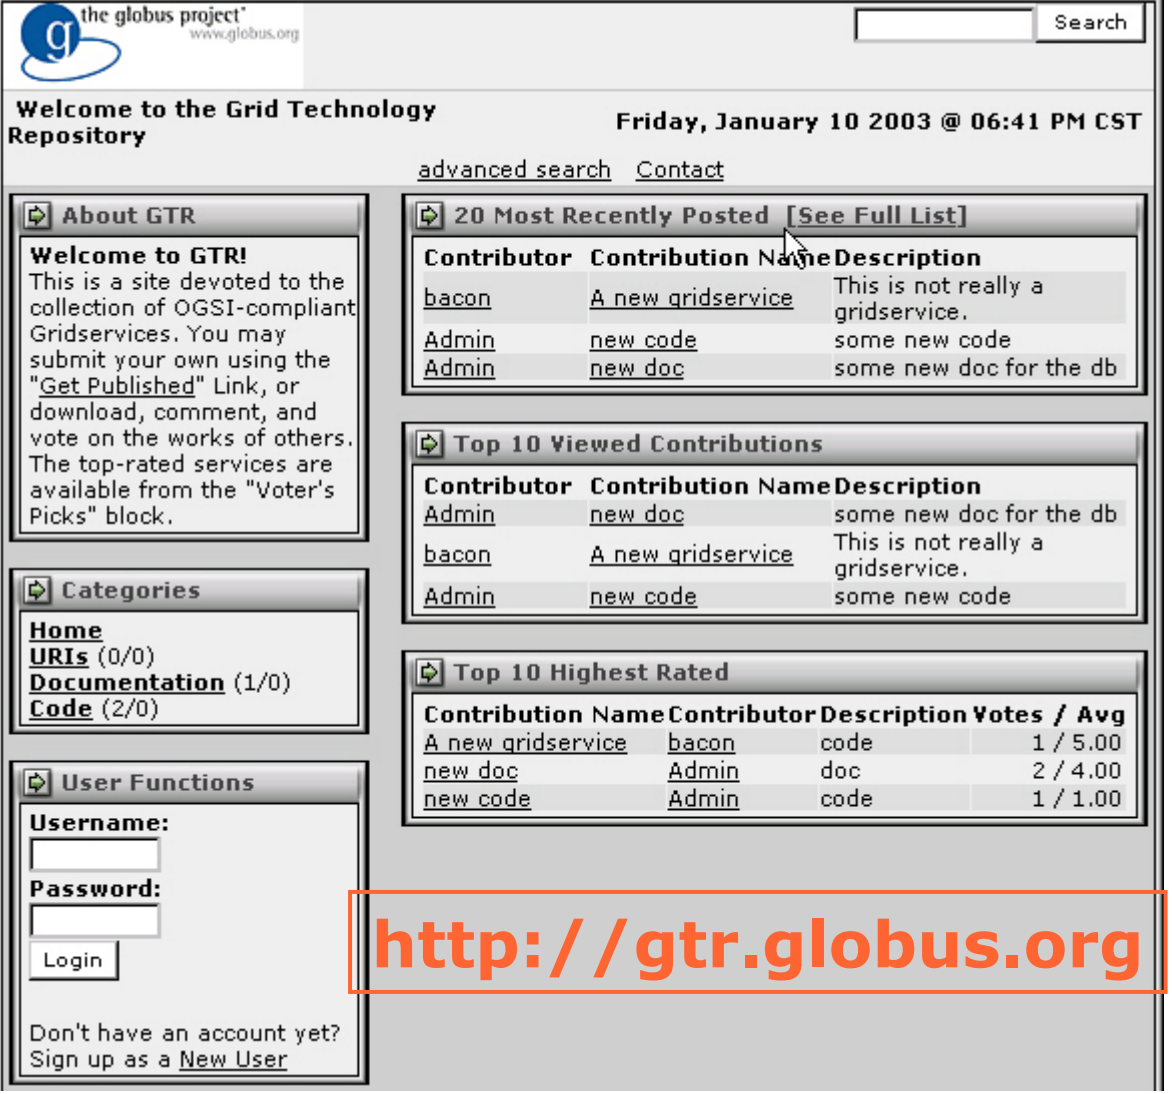

- Community repository
- Clearing house for service definitions, code, documentation
- Encourage collaboration & avoid redundant work

International advisory committee: Ian Foster (Chair), Malcolm Atkinson, John Brooke, Fabrizio Gagliardi, Dennis Gannon, Wolfgang Gentzsch, Andrew Grimshaw, Keith Jackson, Gregor von Laszewski, Satoshi Matsuoka, Jarek Nabrzyski, Bill St. Arnaud, Jay Unger

Overview

- Problem solving in the 21st century
- Open Grid Services Architecture
- Globus Toolkit v3
- **Summary**

The Need for Open Infrastructure

- Broadly deployed services in support of fundamental collaborative activities
  - Formation & operation of virtual organizations
  - Authentication, authorization, discovery, …
- Services, software, and policies enabling on-demand access to critical resources
  - Computers, databases, networks, storage, software services,…
- Operational support for 24x7 availability
- Integration with campus and commercial infrastructures

the globus project™
www.globus.org

Open Infrastructure

Building the National Virtual Collaboratory for Earthquake Engineering Research

NEESgrid

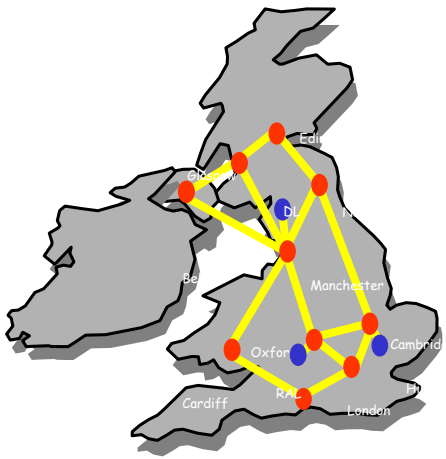

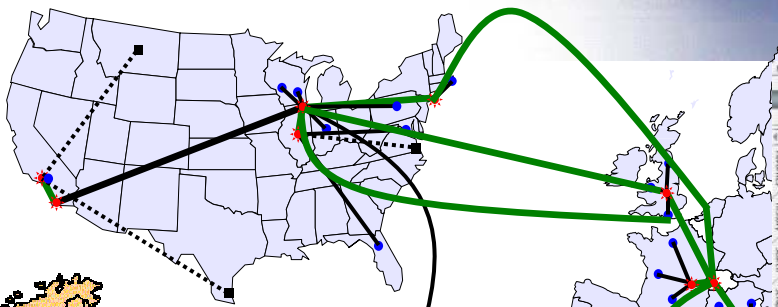

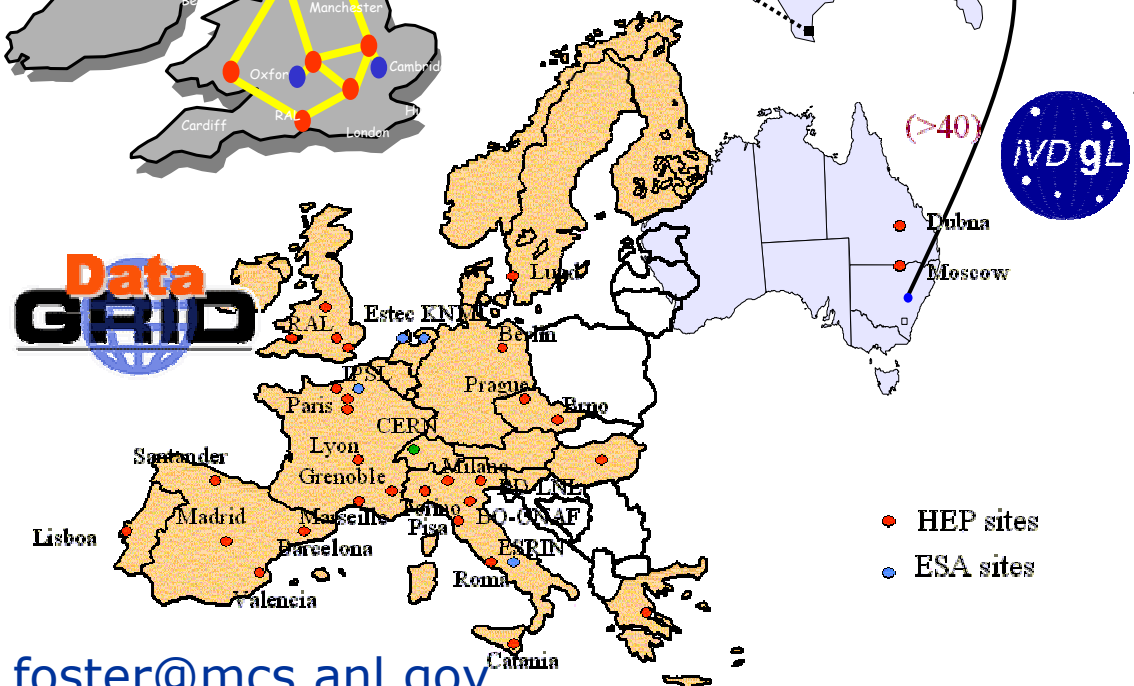

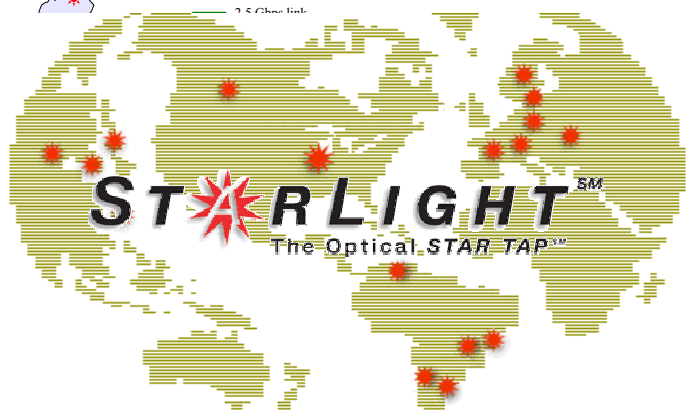

foster@mcs.anl.gov

Summary

- OGSA: standards-based dist. sys. middleware
  - From Web services: standard IDL, discovery, binding independence, other desirable features
  - From Grid/Globus Toolkit: naming, state, lifetime management, etc., etc.
- Rapid progress on definition & implementation
  - OGSI is defined, GT3 implements it (and other things), multiple groups coding to it
  - Much more happening, much more to be done!
- No silver bullet, but a good incremental step towards meeting user requirements

the globus project™
www.globus.org

For More Information

- Open Grid Services Arch.
  - www.ggf.org/ogsa-wg
  - www.globus.org/ogsa
- Global Grid Forum
  - www.ggf.org
- The Globus Project™
  - www.globus.org
- Technical articles
  - www.mcs.anl.gov/~foster