Grid Computing:

Standards and Architecture

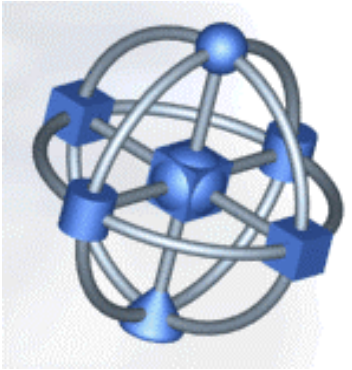

Martin F. Maldonado, Ph.D.

Technical Architect

IBM Grid Computing, Americas

mfmald@us.ibm.com

Contents

Computing Evolution

On Demand
"Dynamic, Responsive, Integrated"

Network-Centric
"The Internet"

Client-Server
"PCs / LANS"

Mainframe
"The Glass House"

e-business

On Demand Operating Environment Attributes

Virtualized

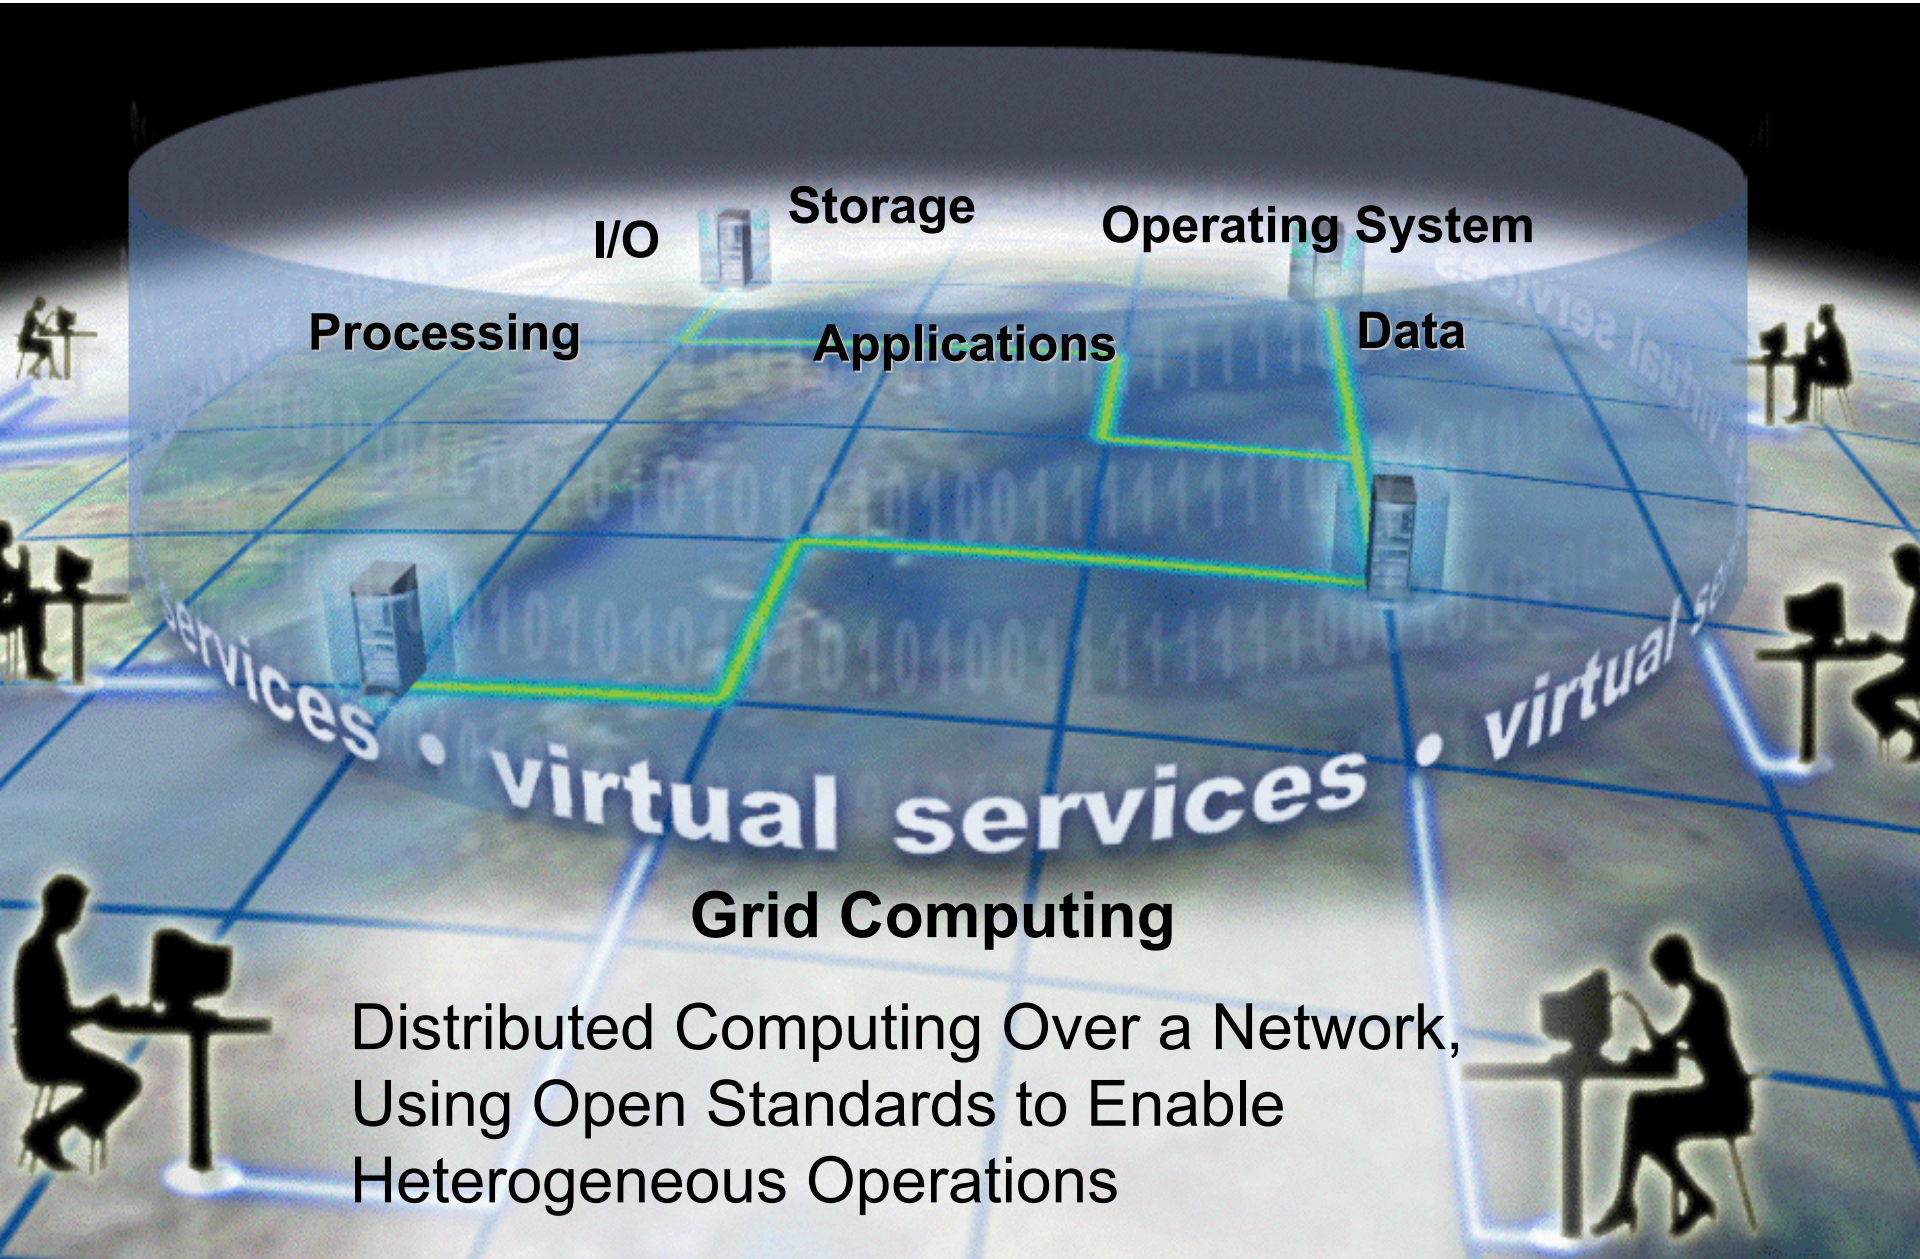

Distributed Computing Over a Network, Using Open Standards to Enable Heterogeneous Operations

What is a Grid?

- **There are three key criteria:**
  - Coordinates resources that are not subject to centralized control …
  - Using standard, open, general-purpose protocols and interfaces …
  - to deliver non-trivial qualities of service.
- **What is not a Grid?**
  - A cluster, a network attached storage device, a scientific instrument, a network, etc.
  - Each is an important component of a Grid, but by itself does not constitute a Grid
  - The web is not (yet) a Grid; its open, general-purpose protocols support access to distributed resource but not the coordinated use of those resources to deliver interesting qualities of service

What is the Grid? A three point checklist, Ian Foster, GRIDToday, July 22, 2002, Vol 1 No. 6

Grid Standards

The Value of Open Standards

Distributed Computing:
Grid
(Globus -> OGSA)

Applications:
Web Services
(SOAP, WSDL, UDDI)

Operating System:
Linux

Information:
World-wide Web
(html, http, j2ee, xml)

Communications:
e-mail
(pop3,SMTP,Mime)

Networking:
The Internet
(TCP/IP)

e-business

Open Grid Services Architecture (OGSA)

"The TCP/IP of Grid Computing"

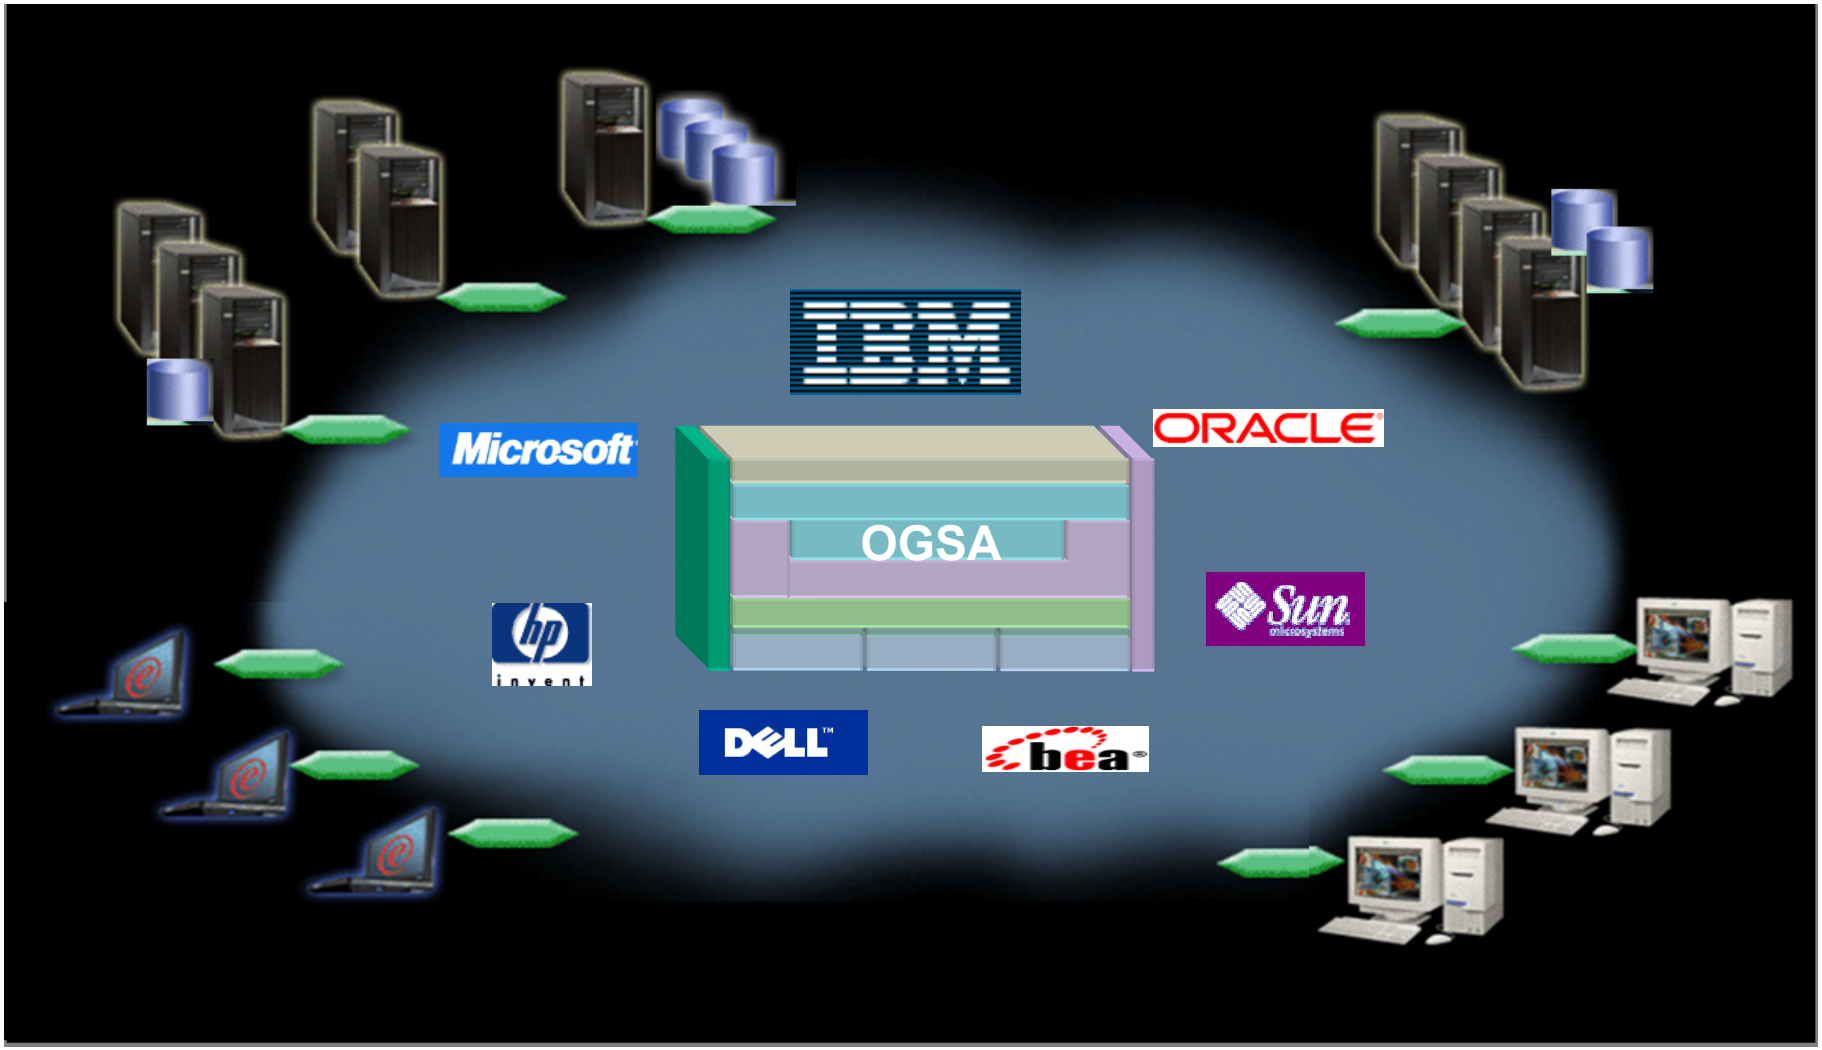

Global Grid Forum

A community-initiated forum of 5000+ individual researchers and practitioners working on distributed computing, or "grid" technologies.

Formed in 2001 by a Merger of Grid Organizations

- European eGrid
- US Grid Forum
- Asia Pacific Grid Community

Primary objective is to promote and support the development, deployment, and implementation of Grid technologies and applications via the creation and documentation of "best practices" - technical specifications, user experiences, an implementation guidelines.

Participants come from over 400 organizations in over 50 countries, with financial and in-kind support coming from sponsor members including technology produce and consumers, as well as academic and federal research institutions.

Modeled After IETF and IRTF

- Meets Three Time Per Year
- Areas, Working Group and Research Groups
- Consensus Based
- Open Membership, Most Work Done on Mailing Lists

IBM is a Platinum Sponsor Member

- Member of Steering Committee
- Member of External Advisory Committee
- Area Directors
- Working Group Chairs

Source: www.ggf.org

GGF Sponsors

Charter Sponsor Members

- Argonne National Laboratory
- NASA Information Power Grid

2002 Platinum Sponsor Members

- Compaq
- Hewlett Packard
- IBM
- Microsoft
- Platform Computing
- Qwest Communications
- Sun Microsystems
- SGI
- US Department of Energy (DOE), Office of Scientific Computing Research
- US National Science Foundation, Division for Advanced Computational Infrastructure and Research (NSF-ACIR)

2002 Gold Sponsor Members

- Level 3 Communications
- Intel
- National Computational Science Alliance (NCSA)
- San Diego Supercomputer Center (SDSC)
- National Institute of Advanced Industrial Science and Technology, Japan (AIST)

2002 Silver Sponsor Members

- Avaki
- Entropia
- Fujitsu America
- Hitachi
- InSORS Integrated Communications
- Johnson & Johnson
- United Devices
- University of Virginia

42 GGF Groups as of January 2003

Applications and Programming Environments

Grid Checkpoint/Recovery
Advanced Programming Models (APM-RG)
Grid Computing Environments (GCE-RG)
Life Sciences Grid RG
Advanced Collaborative Environments (ACE-RG)
Applications and Test Beds (APPS-RG)
Grid User Services (GUS-RG)

Architecture

Open Grid Services Infrastructure (OGSI-WG)
New Productivity Initiative (NPI-WG)
Accounting Models (ACCT-RG)
Service Management Frameworks (JINI-RG)
Open Source Software (OSS-WG)
Open Grid Services Architecture (OGSA-WG)
Grid Protocol Architecture (GPA-RG)
Production Grid Management RG

Data

GridFTP-WG
Data Replication (REPL-RG)
Grid High-Performance Networking (GHPN-RG)
Data Access and Integration Services (DAIS-WG)
Persistent Archives (PA-RG)
Data Transport (DT-WG)

Information Systems and Performance

Discovery and Monitoring Event Description (DAMED-WG)
Grid Information Retrieval (GIR-WG)
Relational Grid Information Services (RGIS-RG)
Semantic Grid RG
Network Measurement (NM-WG)
CIM based Grid Schema (CGS-WG)
Grid Benchmarking (GB-RG)

Peer-to-Peer

Appliance Aggregation
OGSA-P2P-Security

Scheduling and Resource Management

Scheduling Attributes (SA-WG)
Distributed Resource Management Application API (DRMAA-WG)
OGSA Resource Usage Service (RUS-WG)
Usage Record (UR-WG)
Scheduling Dictionary (SD-WG)
Grid Resource Allocation Agreement Protocol (GRAAP-WG)
Grid Economic Services Architecture (GESA-WG)

Security

Grid Security Infrastructure (GSI-WG)
Open Grid Service Architecture Security (OGSA-SEC-RG)
Large Site AAA (AAA-WG)
Grid Certificate Policy (GCP-WG)
CA Ops (CAO-WG)

Source: www.ggf.org

BM Active Industry Participation in GGF

- APE
  - Boeing
- ARCH
  - Avaki, Fujitsu, IBM, Platform, Sun (JINI only)
- DATA
  - Avaki, IBM
- GIS-PERF
  - Platform, IBM
- SCHED
  - IBM, Intel, Sun
- GS
  - IBM, Verisign

Open Grid Services Architecture

Open Grid Services Architecture Objectives

- **Distributed Resource Management across heterogeneous platforms**
- **Seamless QoS delivery**
- **Common Base for Autonomic Management Solutions**
- **Common infrastructure building blocks to avoid "stovepipe solution towers"**
- **Open and Published Interfaces**
- **Industry-standard integration technologies**
  - web services, soap, xml...
- **Seamless integration with existing IT resources**
  - Separate interface from implementation

Distributed Computing: A Common Problem

- **Web services, Autonomic computing and Grid efforts all try to address aspects of distributed computing:**
  - Defining an open distributed computing paradigm.
  - Dealing with heterogeneous platforms, protocols and applications.
- **GRID has focused on Scientific / Technical Computing across organizational boundaries**
  - Here, secure, distributed Resource Sharing is the key
  - But no standards exist for inter-operability or pluggable components
- **Web Services initial focus has been on application integration**
  - not resource provisioning or system integration
- **Autonomic computing is focused on managing commercial IT infrastructures:**
  - Here, sharing resources is not the issue: Managing them is!
  - Sharing function is not the issue: Building solutions on top is!

The Best of Both Worlds

Web Services & Grid Protocols

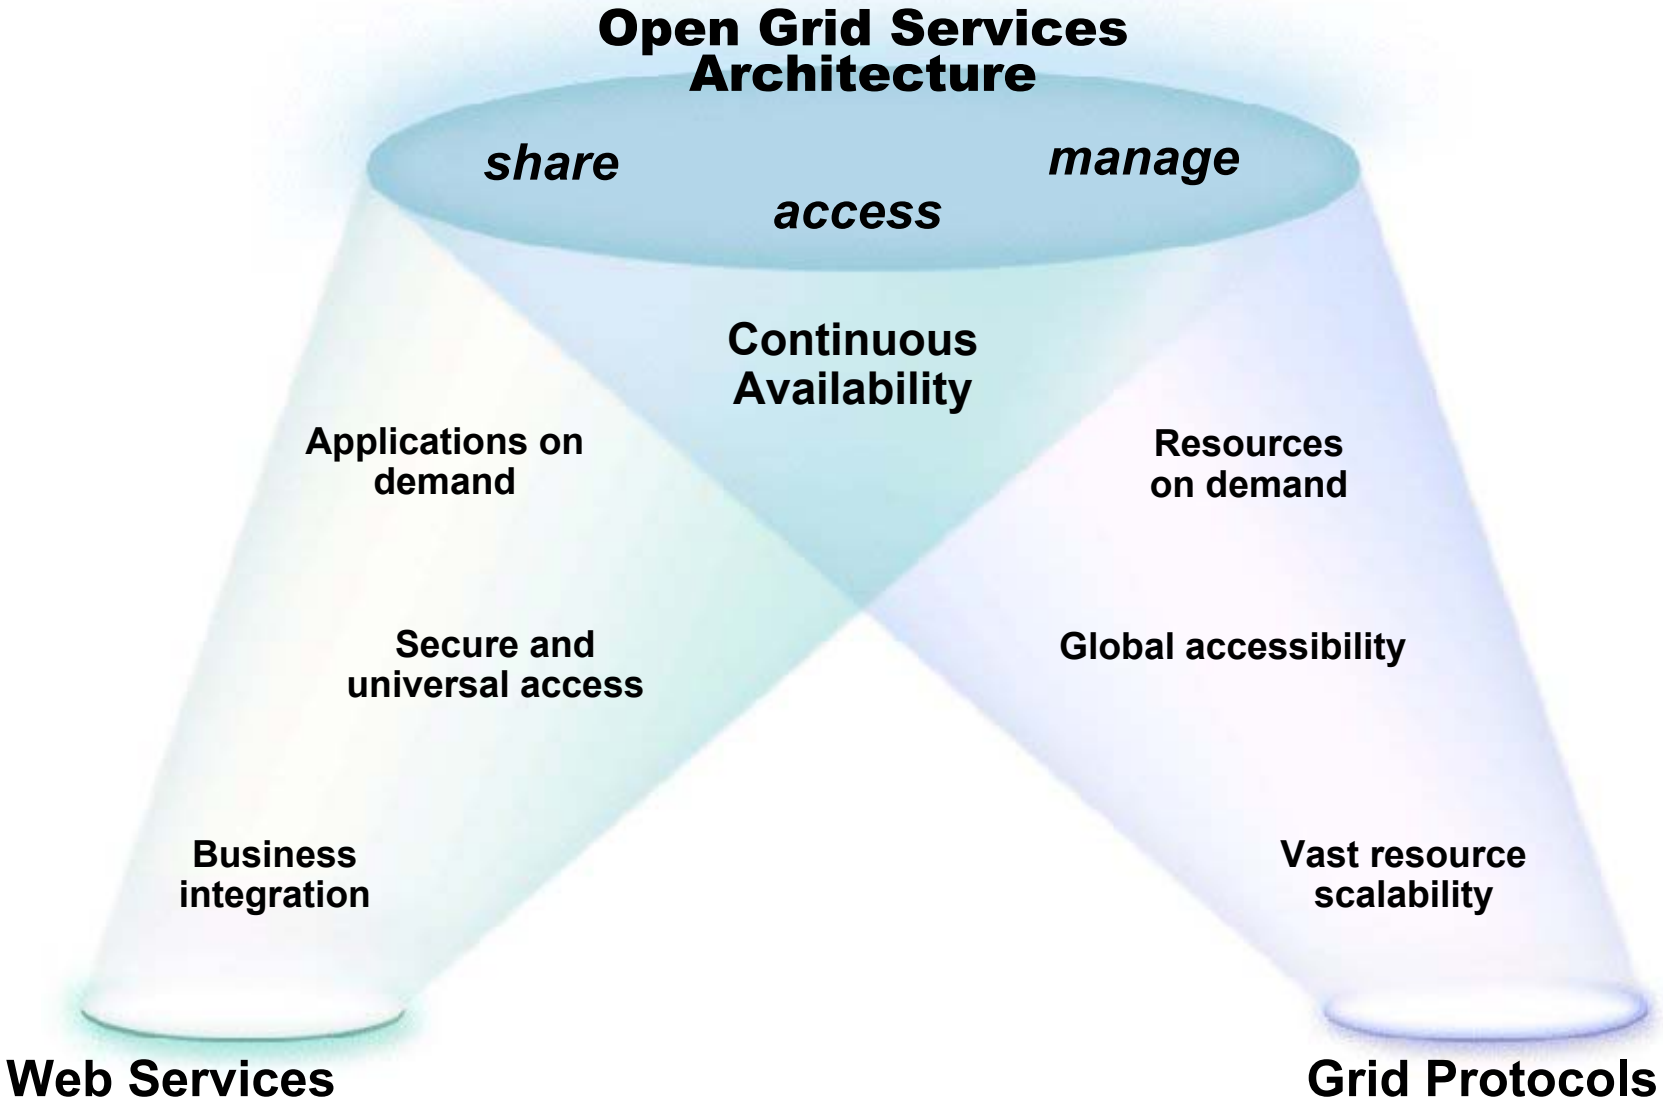

Architecture Framework

OGSA Structure

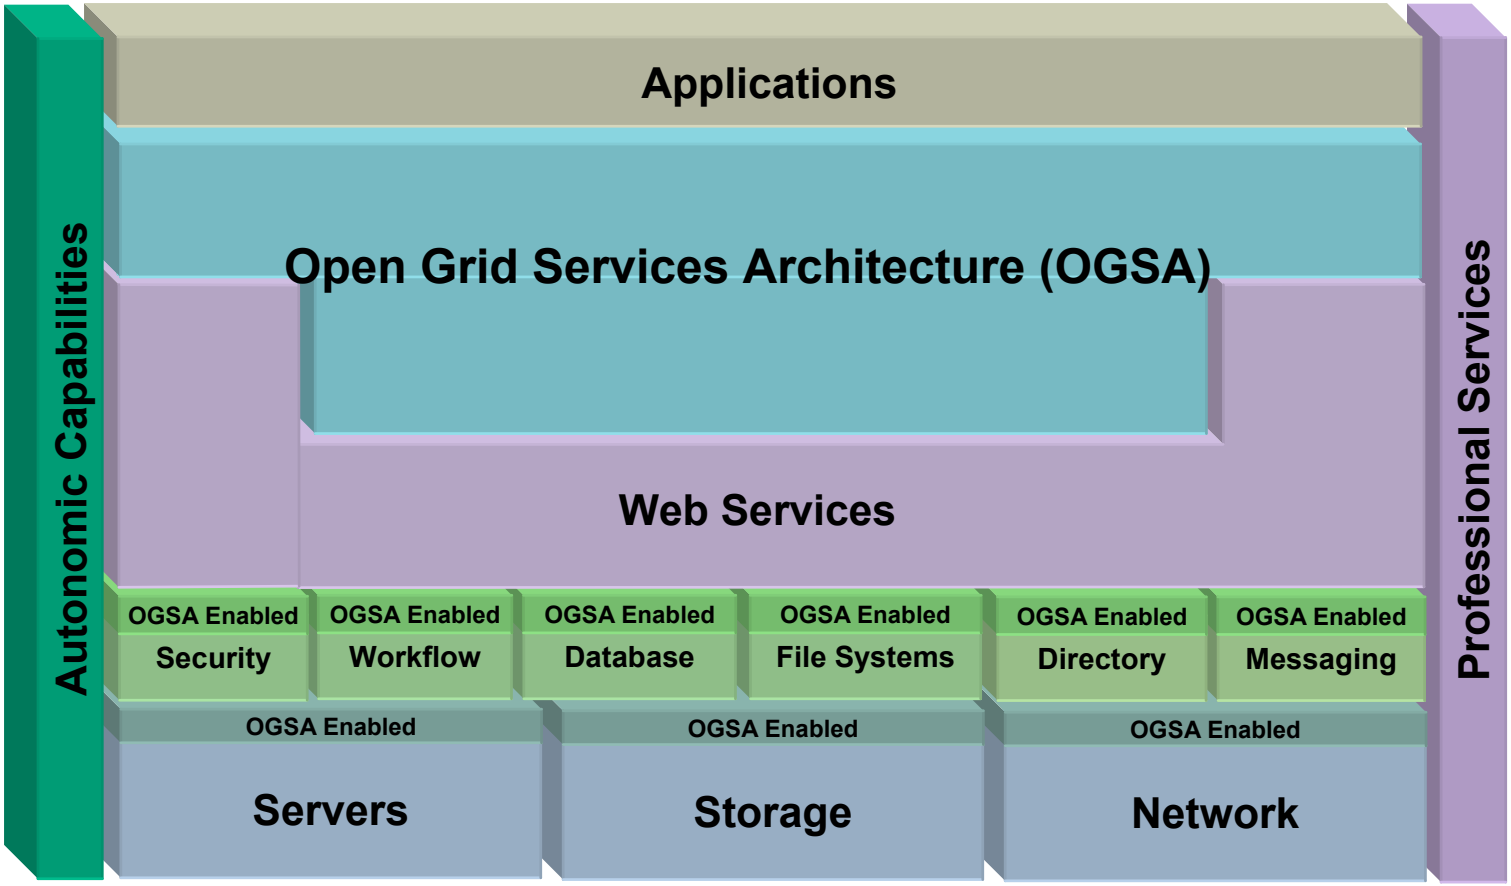

Architecture Framework

OGSA Structure

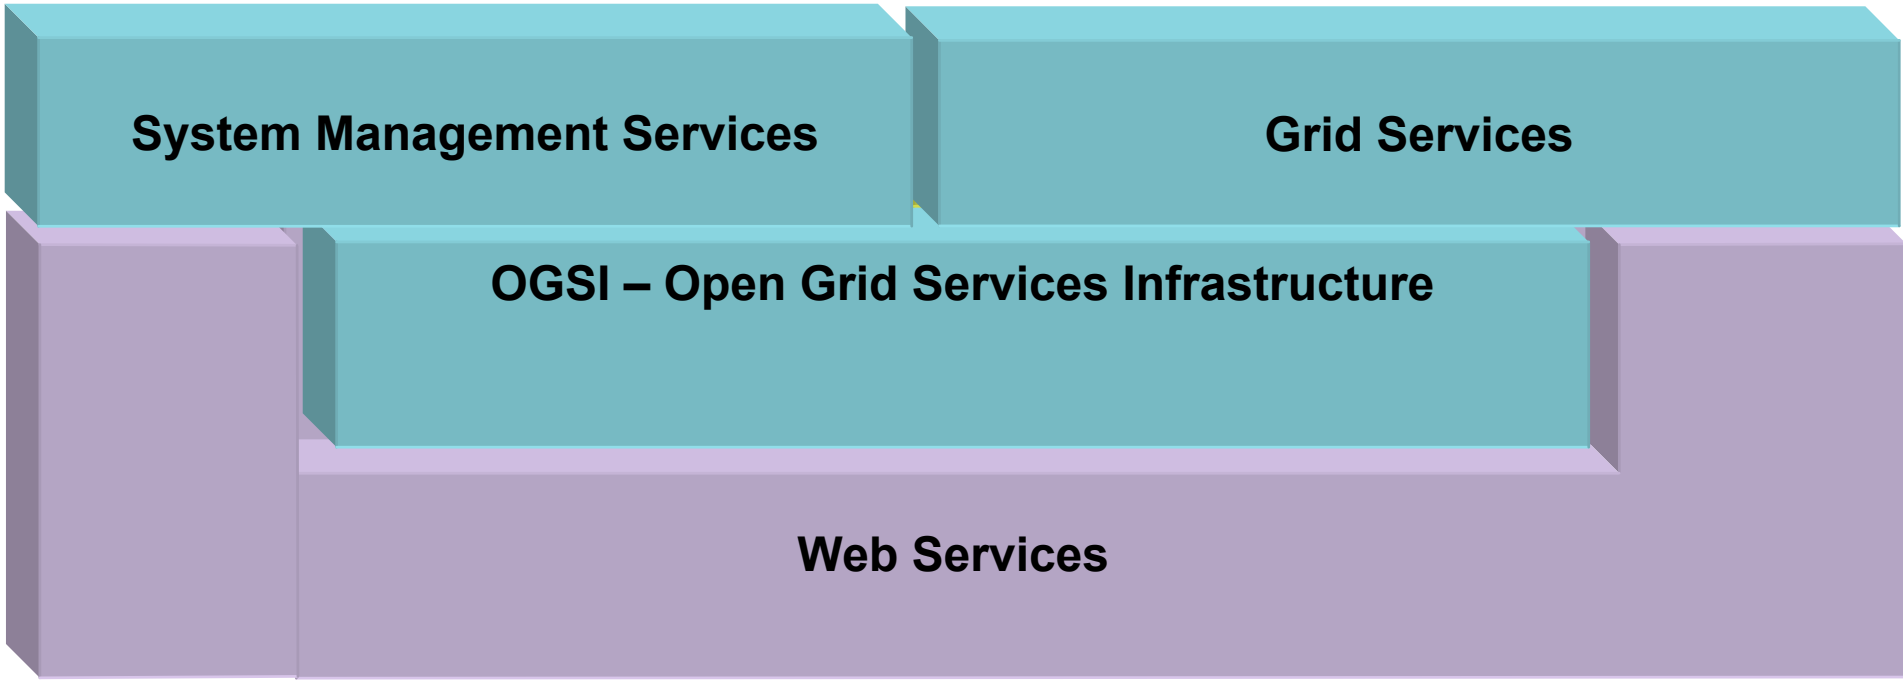

Architecture Framework

OGSA Structure – OGSI

- **Exploits existing web services properties**
  - Interface abstraction (WSDL)
  - Protocol, language, hosting platform independence
- **Enhancement to web services**
  - State Management
  - Event Notification
  - Referenceable Handles
  - Lifecycle Management
  - Service Data Extension

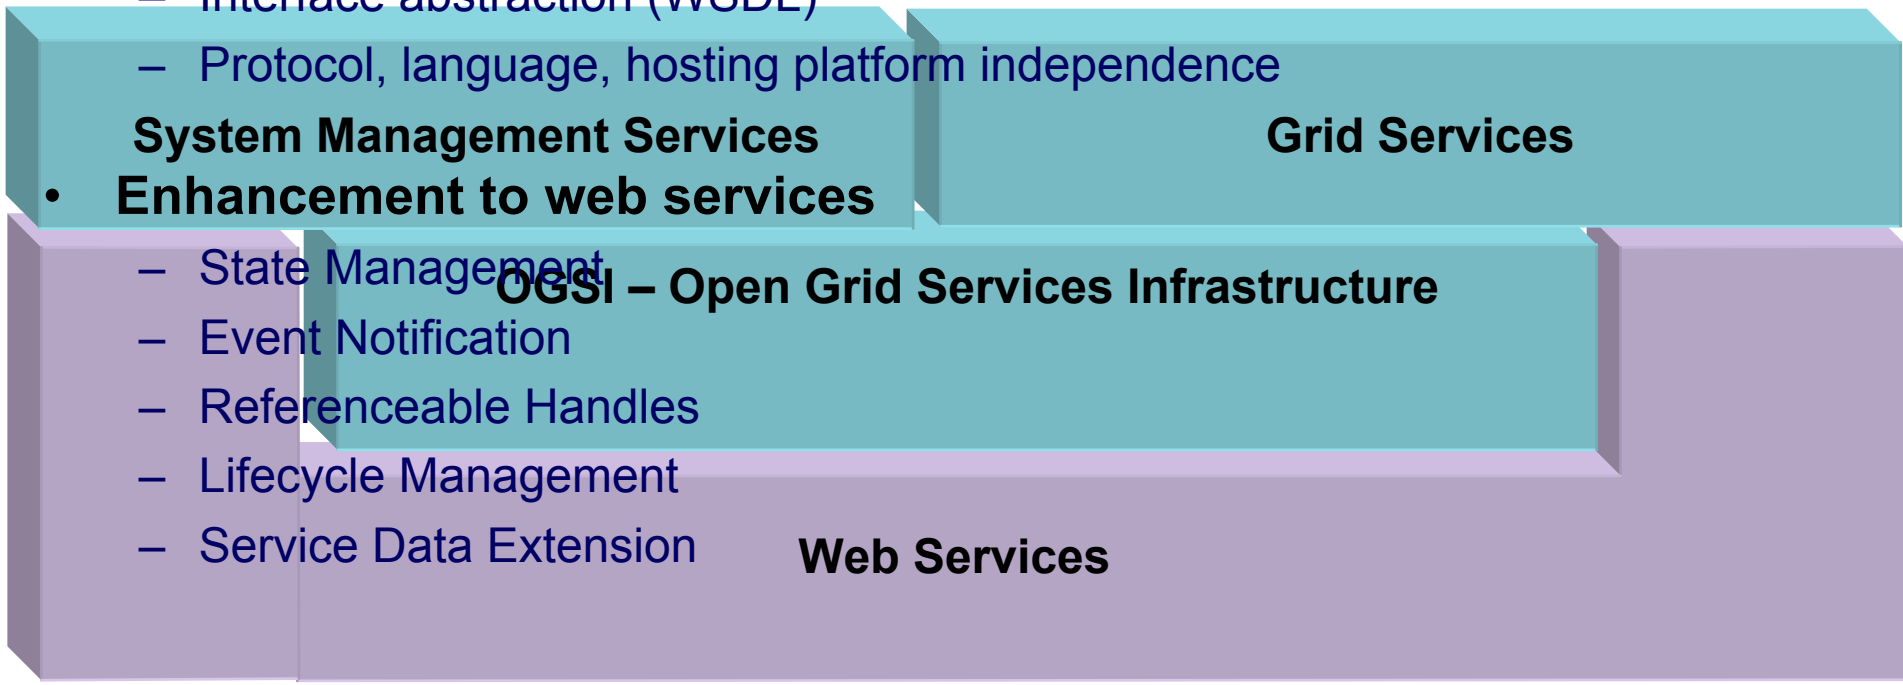

Architecture Framework

OGSA Structure

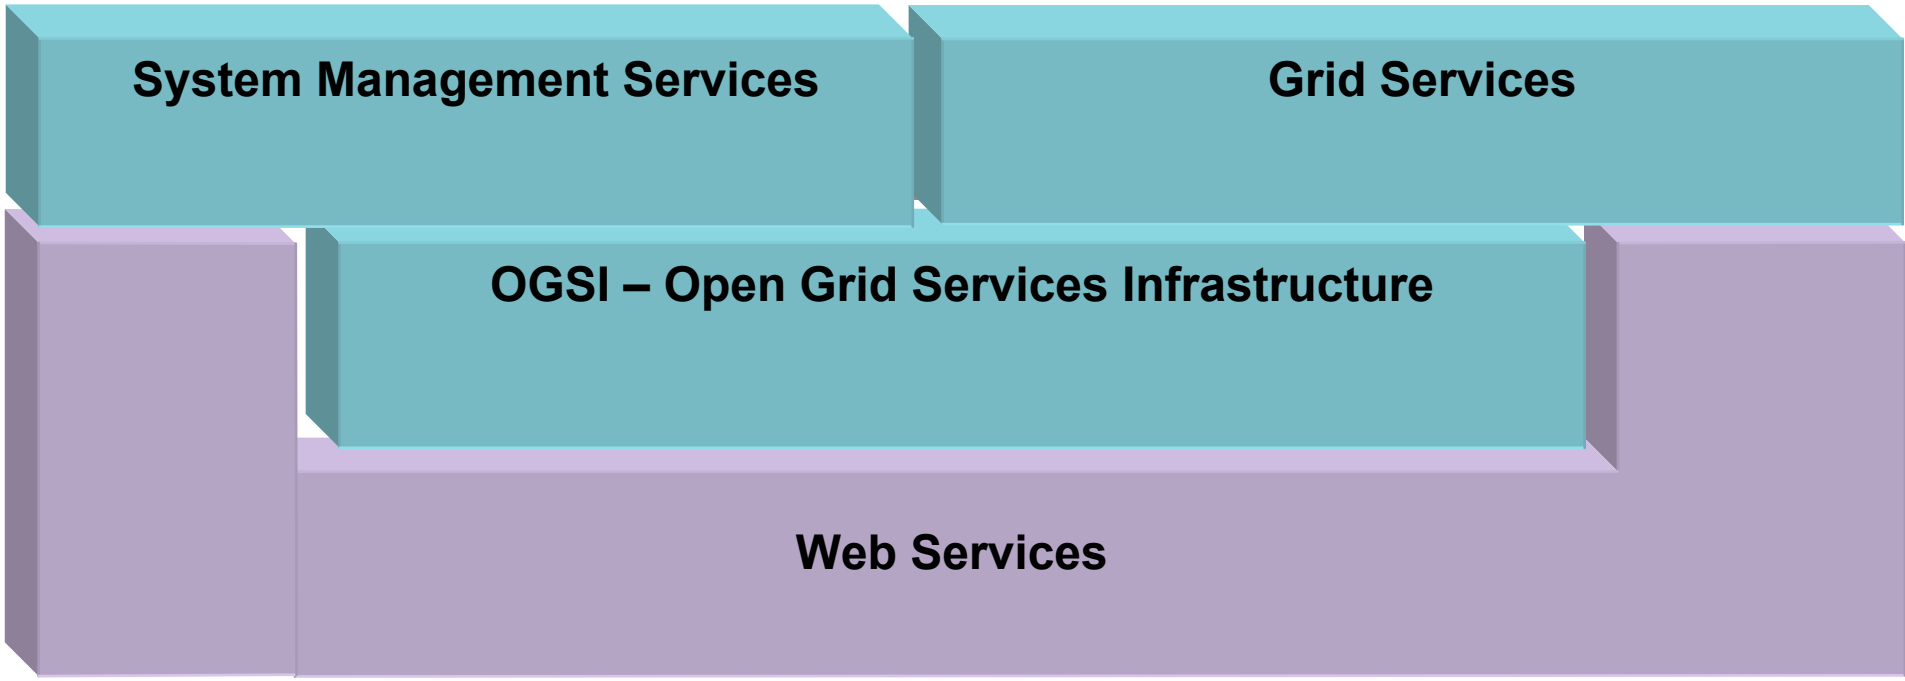

Architecture Framework

Products and Services for Grids

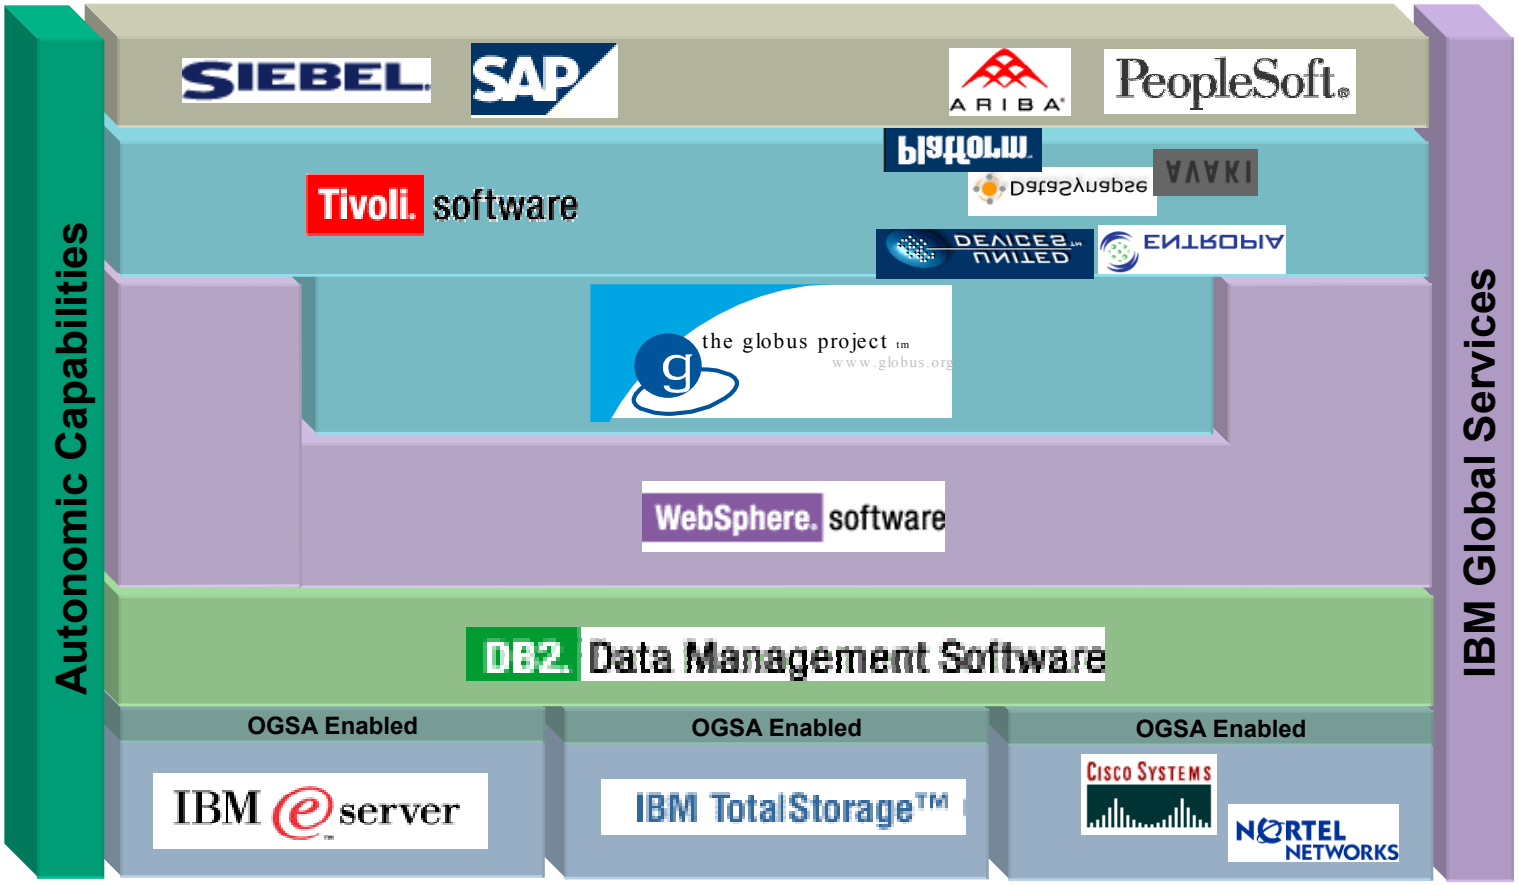

Grid Services

OGSA Services Model

- Everything is represented by a (Grid) service
- A service is a network-enabled entity that provides some capability
- A service can be a computation resource, storage resource, network, program, database, and so on
- Services can be transient, created dynamically and destroyed when no longer needed
- Separates the definition of the interface and protocols to invoke the interface
- Simplifies virtualization - encapsulation behind a common interface of diverse implementations
- Virtualization allows:
  - consistent resource access across multiple heterogeneous platforms with local and remote transparency
  - enable mapping of multiple logical resource instances onto the same physical resource
  - management of resources based on composition from lower-level resources
  - allows the composition of services to form more sophisticated services

Hosting Environment

- OGSA does not address issues of implementation programming model, programming language, implementation tools, or execution environment
- Grid services are instantiated within a specific hosting environment
- Host environment defines how a Grid service meets it obligation to Grid service semantics
  - rely on native operating system processes, implementing service in a variety of languages
  - implemented on container or component-based hosting environment such as J2EE, Websphere, .NET, and Sun One

Open Grid Infrastructure (OGSI)

Anatomy of a Grid Service

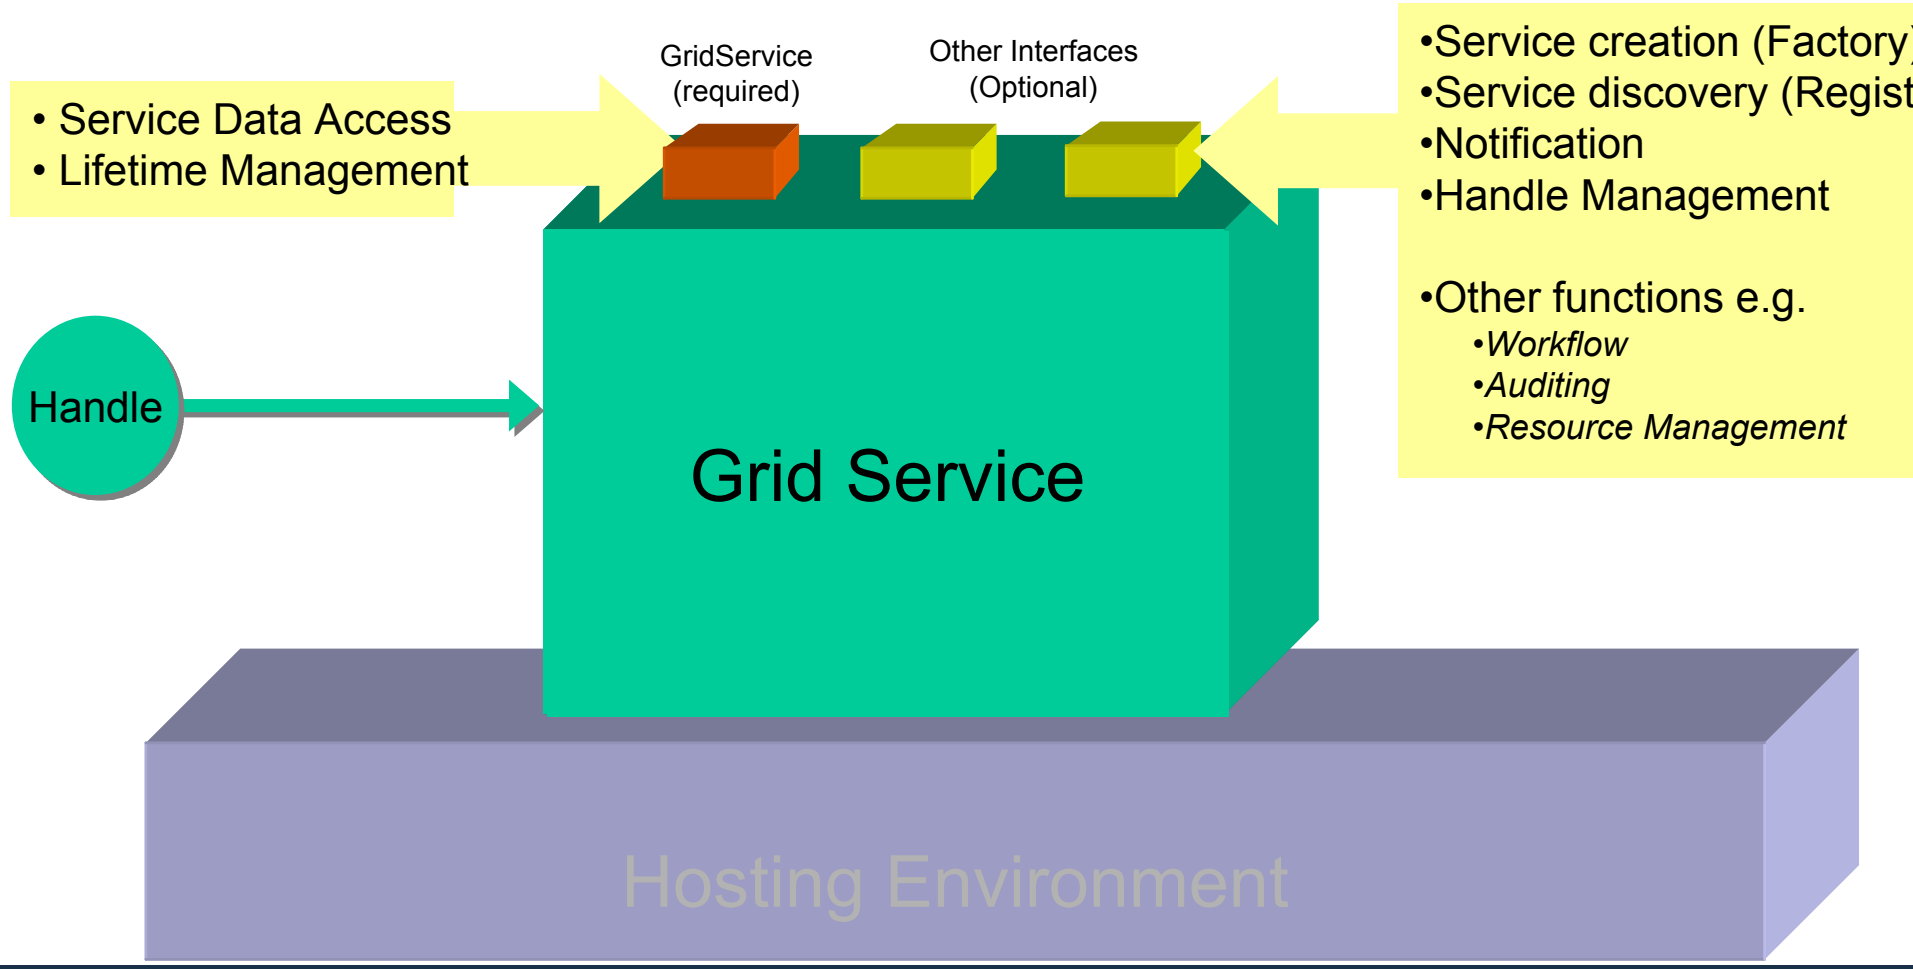

Open Grid Infrastructure (OGSI)

Grid Service Implementation Independence

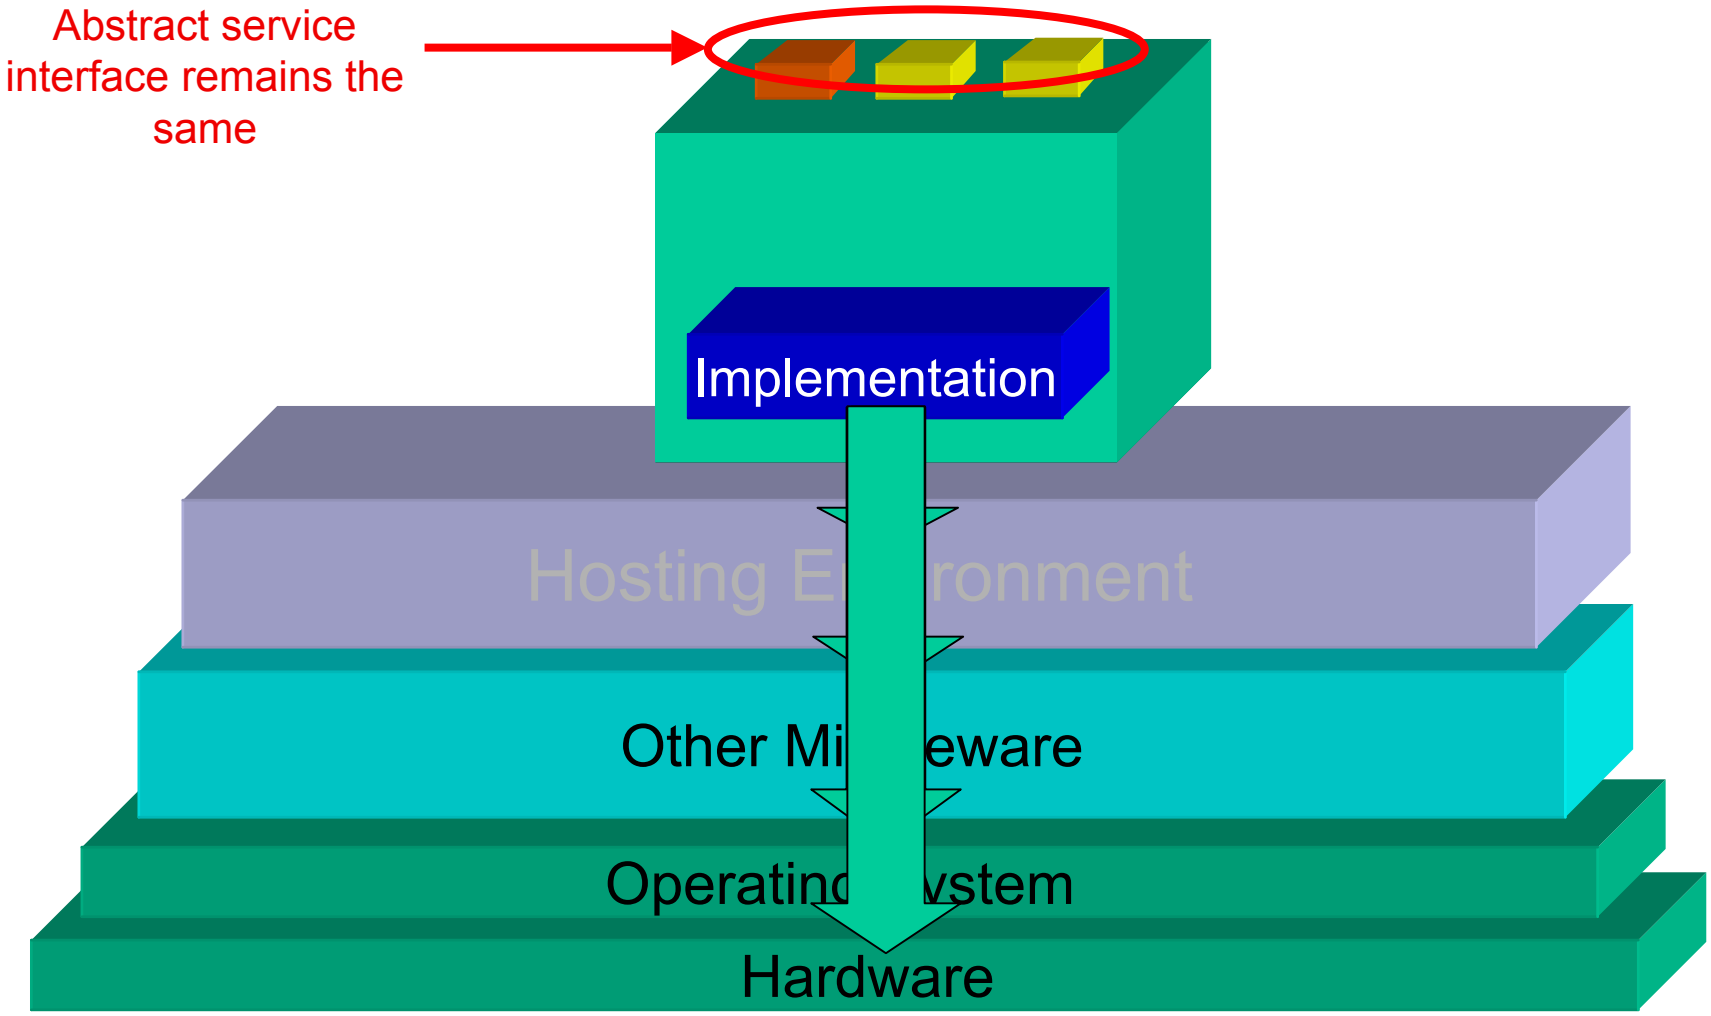

Open Grid Infrastructure (OGSI)

Grid Service Implementation - Examples

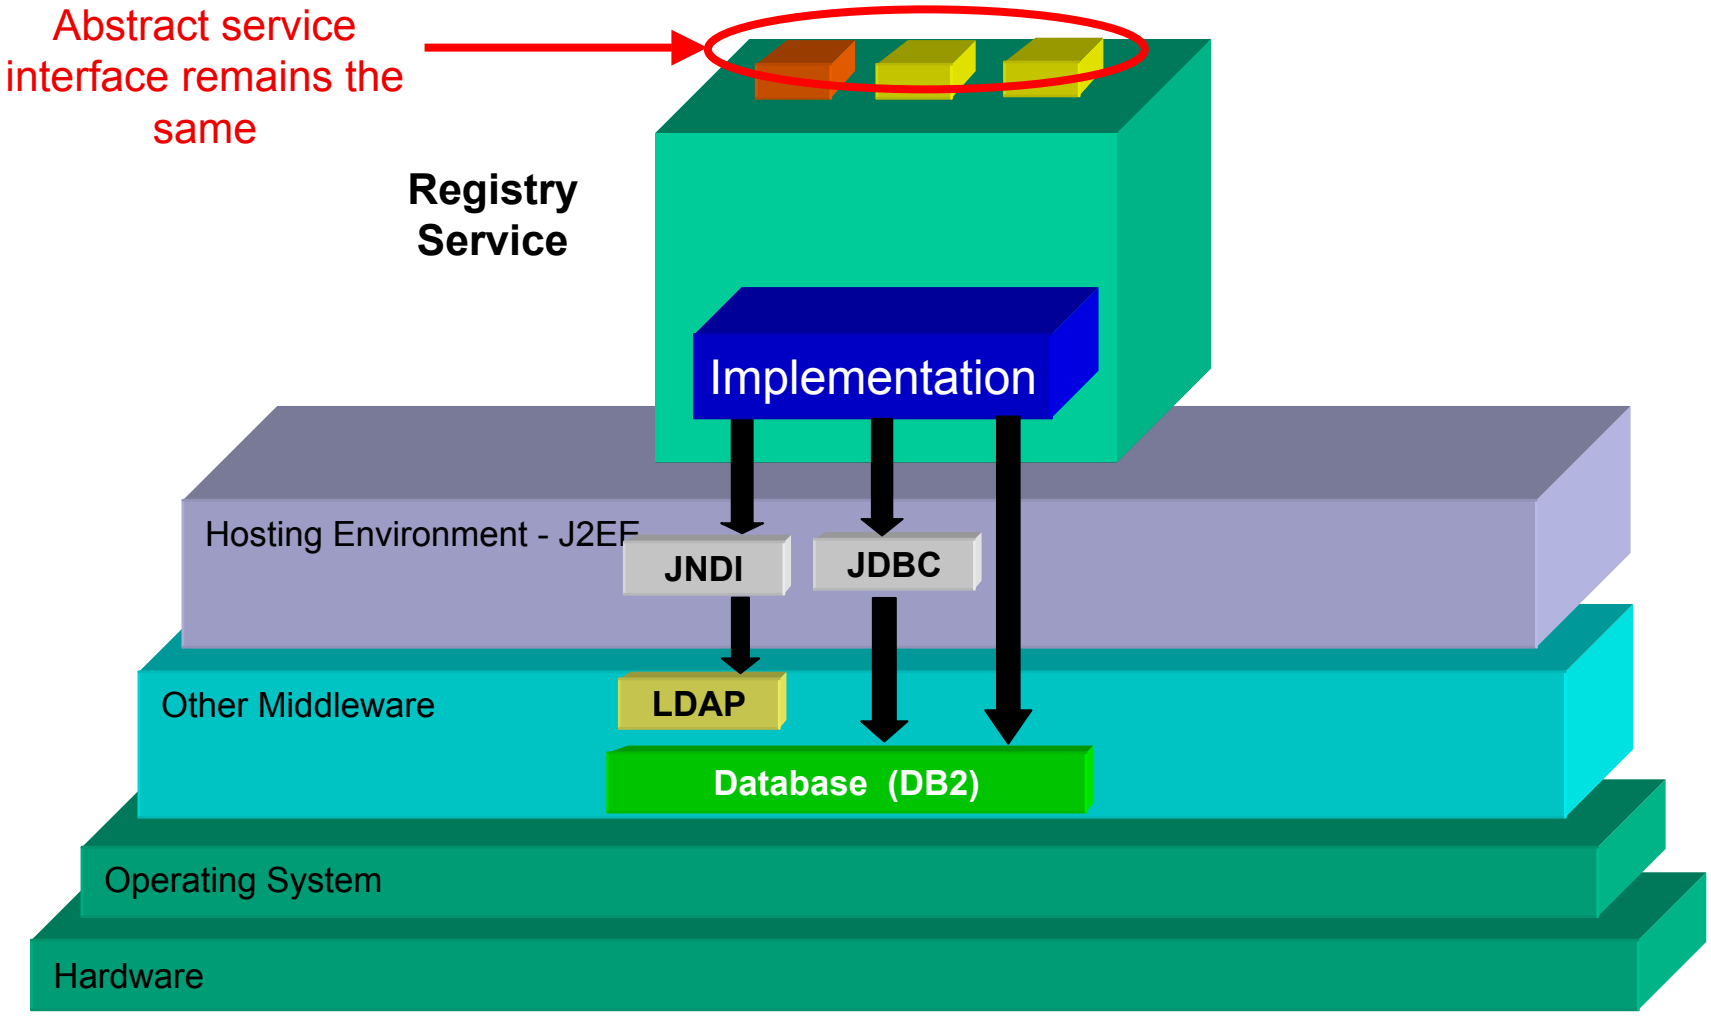

Open Grid Infrastructure (OGSI)

Grid Service Implementation - Examples

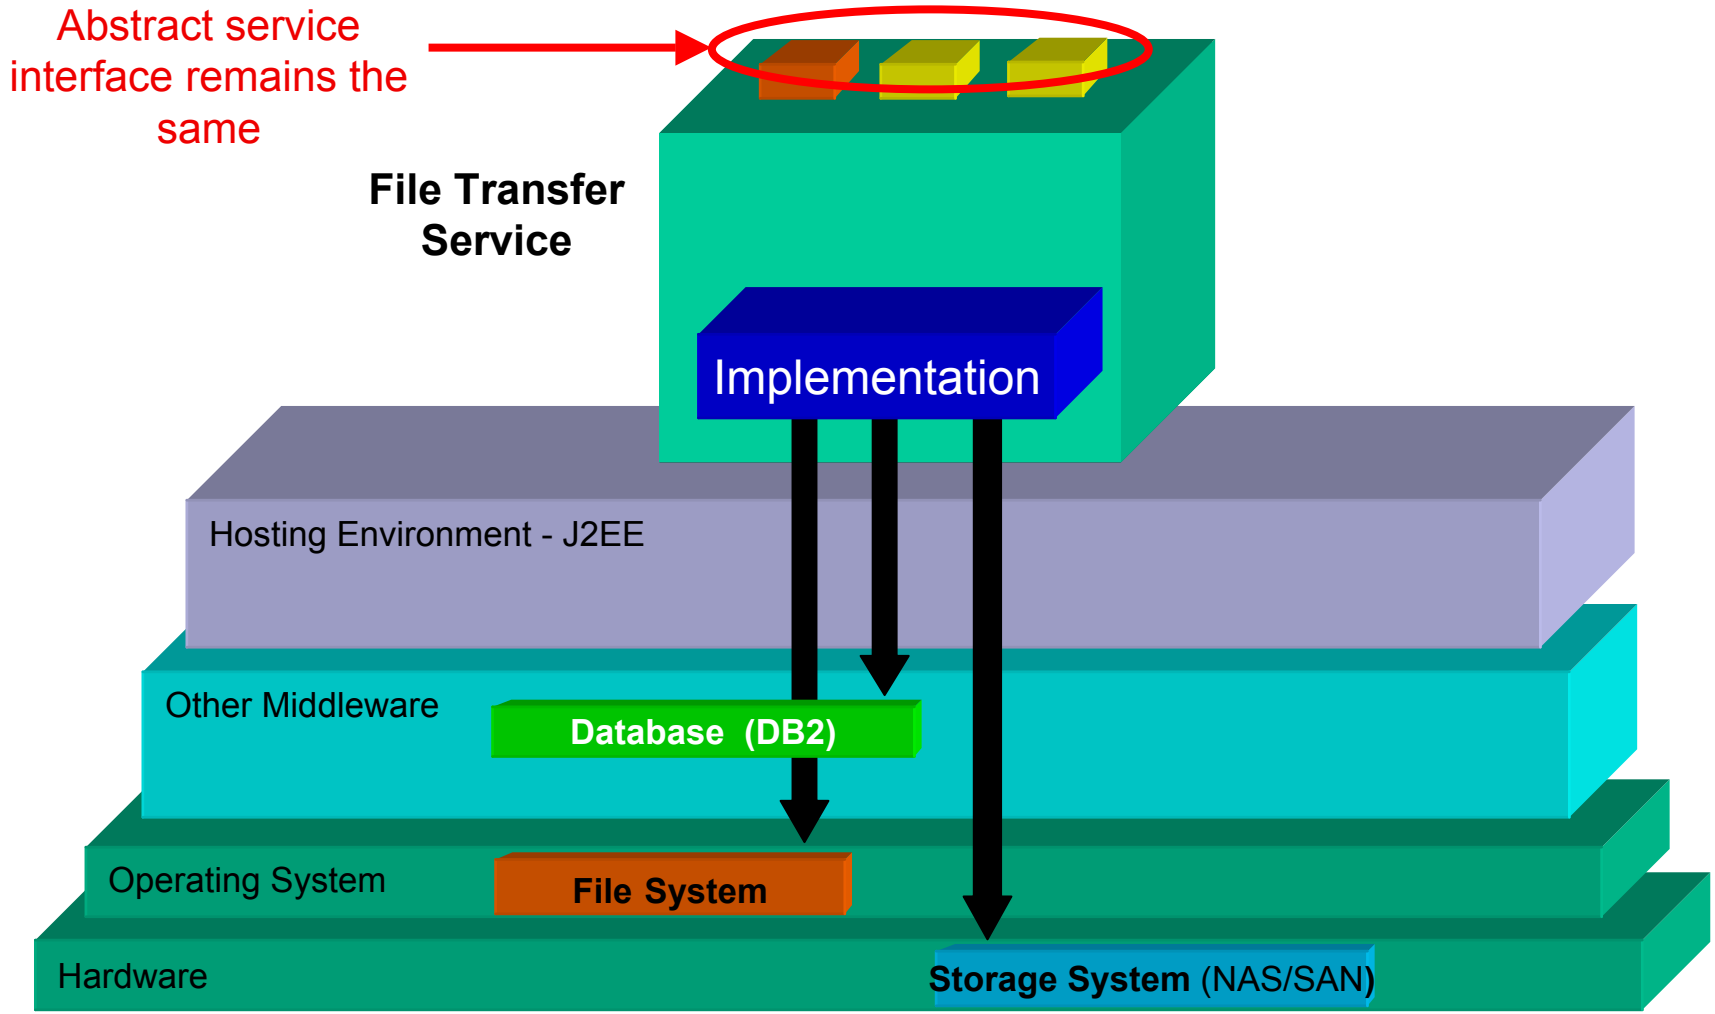

Grid Data Access and Integration

Architectural Principals

- **Heterogencity Transparency**
  - The access mechanism should be independent of the actual implementation
- **Location Transparency**
  - An application should be able to access data irrespective of its location
- **Name Transparency**
  - An application should be able to access data without knowing its name or location
  - Data access should be via logical domains, qualified by predicates on attributes of the desired object
- **Distribution Transparency**
  - An application should be able to query and update data without being aware that it comes from a set of distributed sources
- **Replication Transparency**
  - Grid data may be replicated or cached in many places for performance and availability
- **Ownership and Costing Transparency**
  - Applications should be spared from separately negotiating for access to individual sources, whether in terms of access authorization, or in terms of access costs.

Source: Grid Database Access and Integration: Requirements and Functionalities

Principal portTypes

- **GridDataService**
  - Service Data Elements
    - Logical Schema
    - Physical Schema
    - StatementNotificationTypes
    - ResultFormatTypes
    - DatabaseTypes
    - SystemName
    - TransactionCapability
    - preparedStatements
    - resultCollections
  - Operation
    - perform
  - Messages
    - gridDataServiceRequest
    - gridDataServiceResponse
- **GridDataTransport**
  - Service Data Elements
    - LogicallySupportedTypes
    - PhysicallySupportedTypes
    - activeBlocks
  - Operations
    - perform
  - Messages
    - GridDataTransportStatement
    - GridDataTransportResponse
    - GridDataTransportFault

Source: Grid Database Service Specification

Creating and Using Grid Data Services

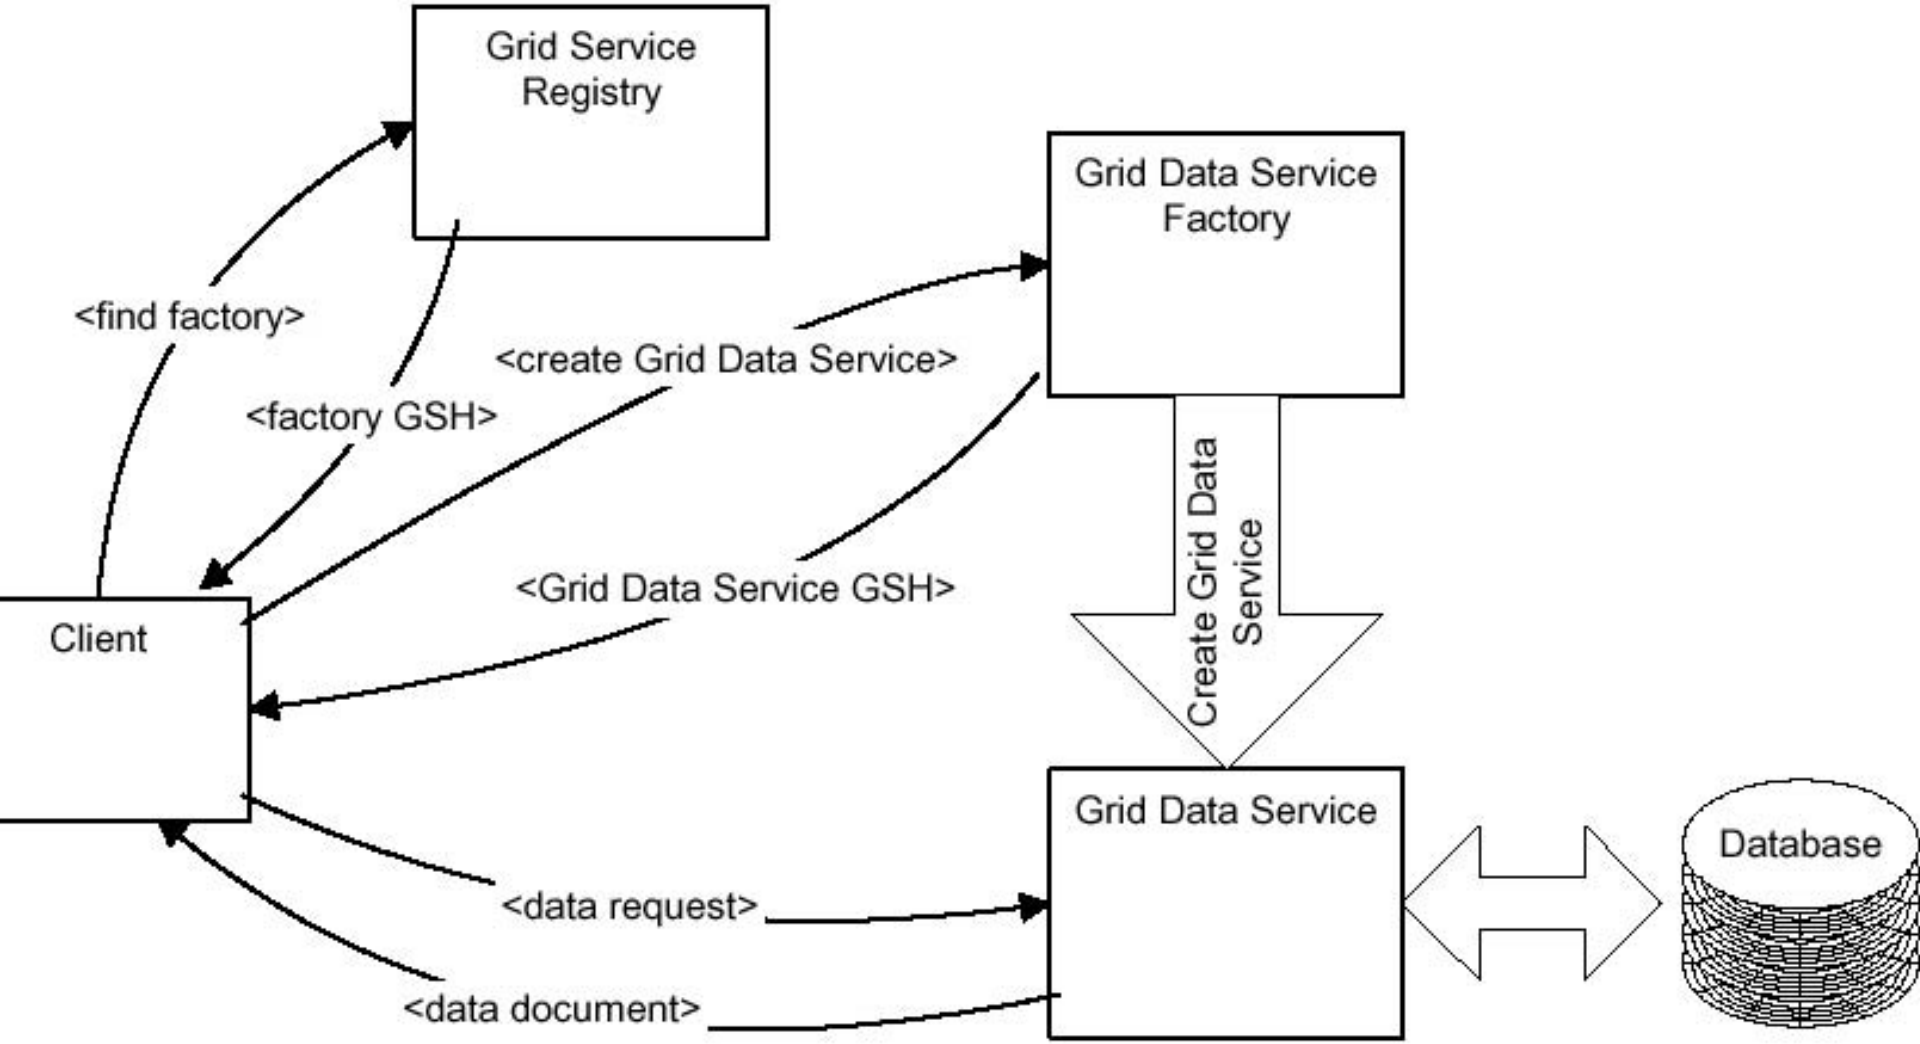

Source: Grid Database Service Specification

Requestor Retrieving Data from Grid Data Servic

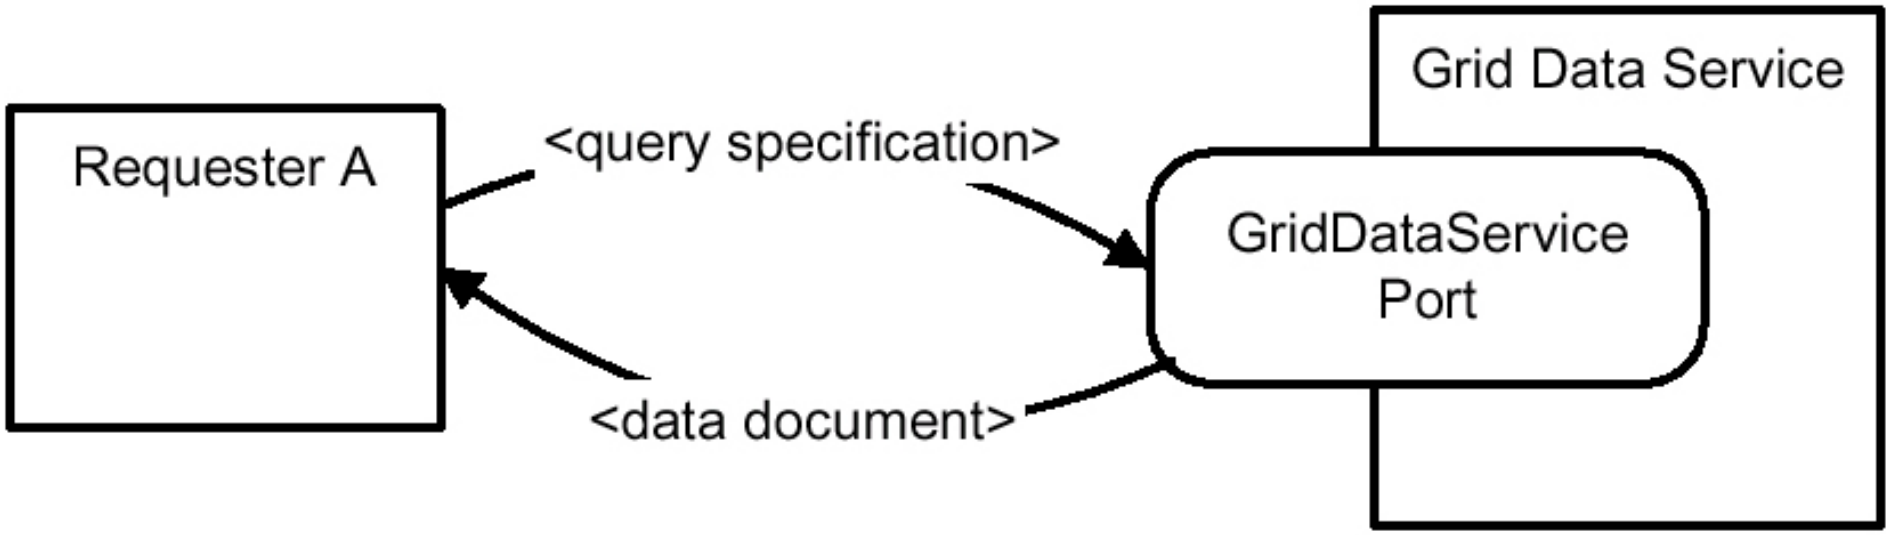

Source: Grid Database Service Specification

Requestor Using Grid Services Ports

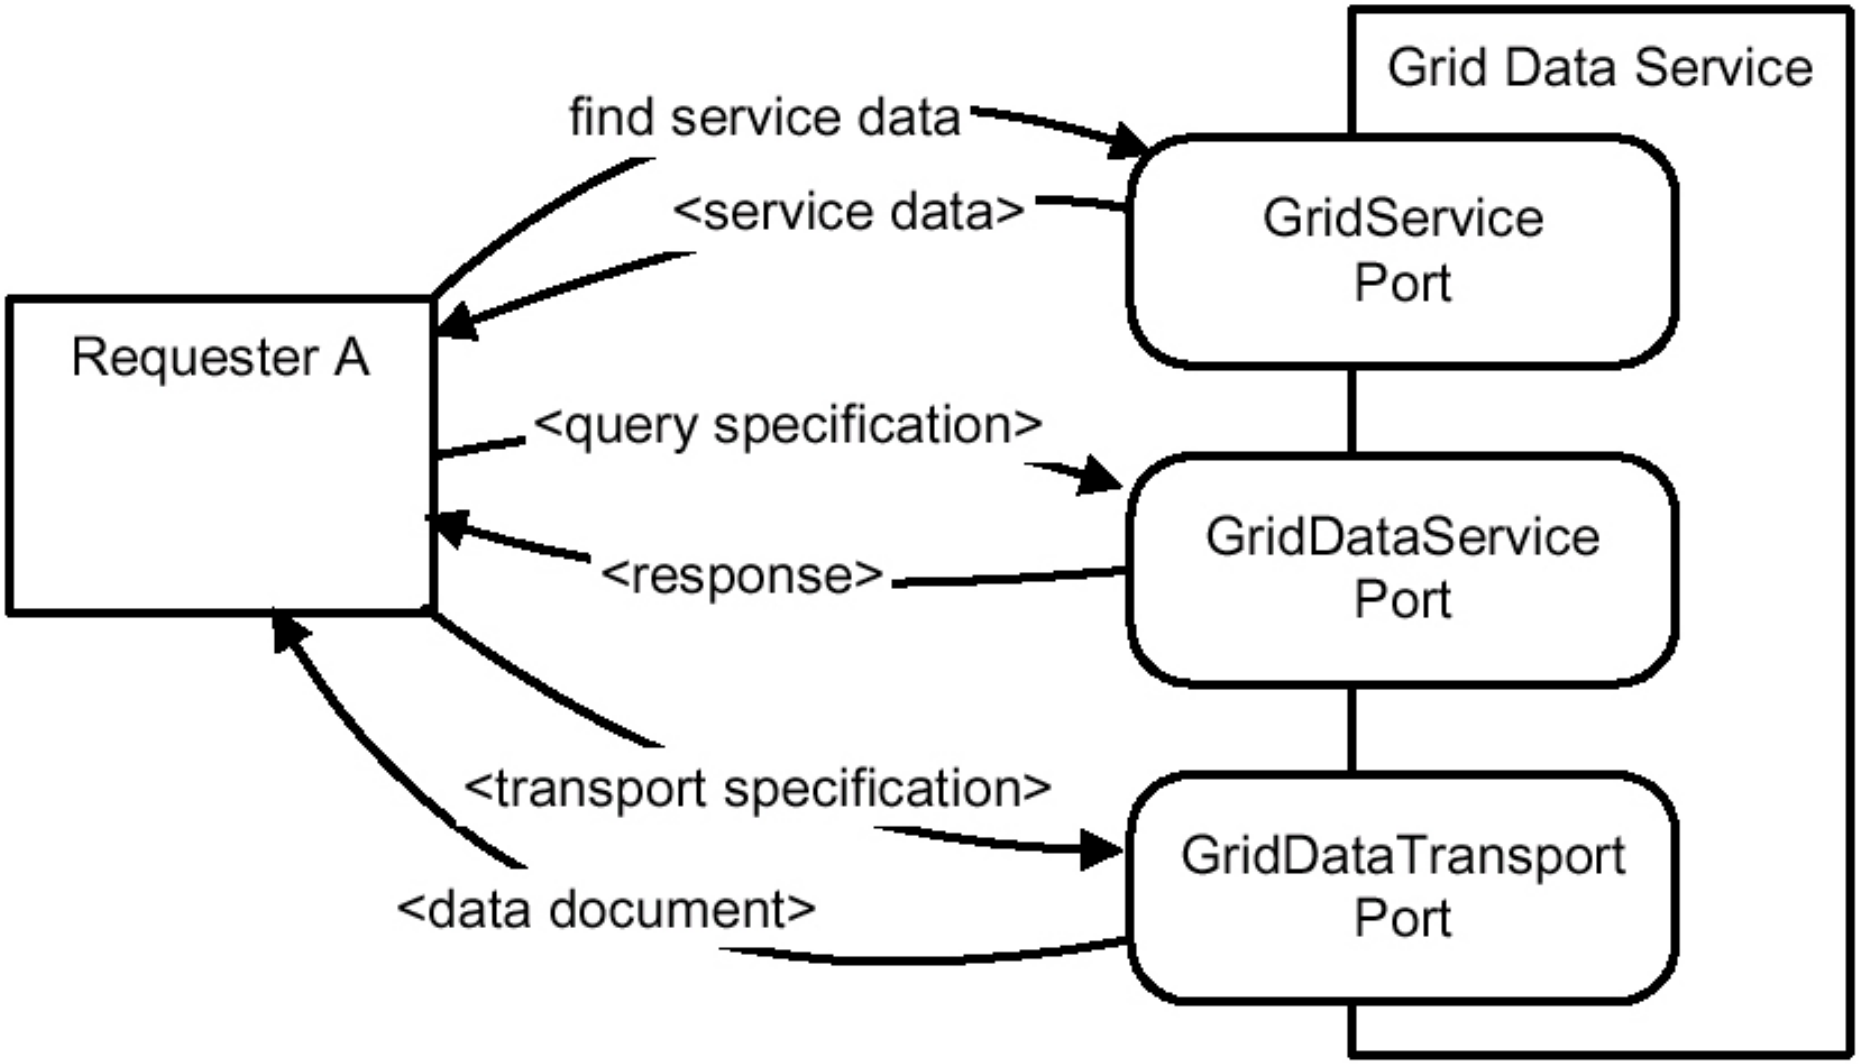

Source: Grid Database Service Specification

Query Request with Deliver to Third Parties

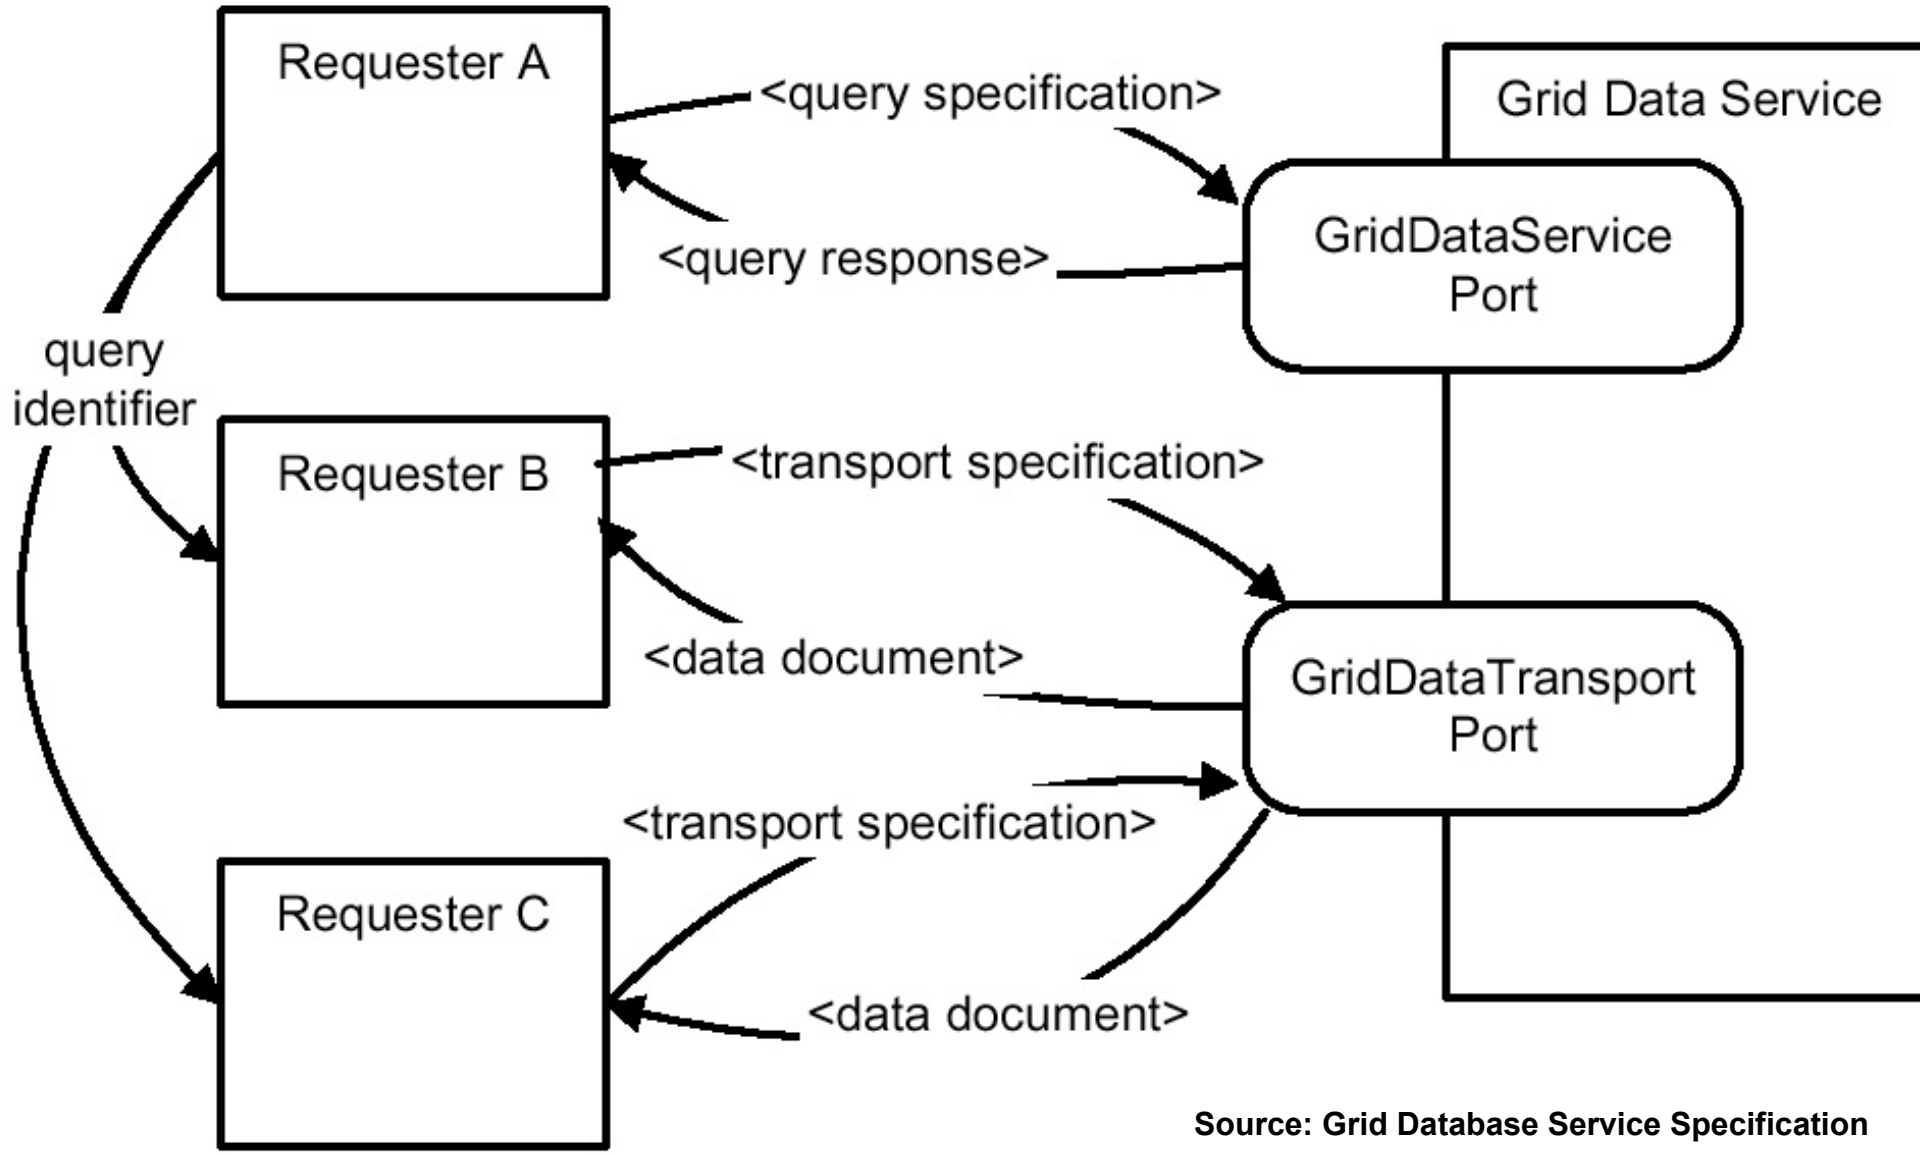

Source: Grid Database Service Specification

Sending Data from one GDS to Another

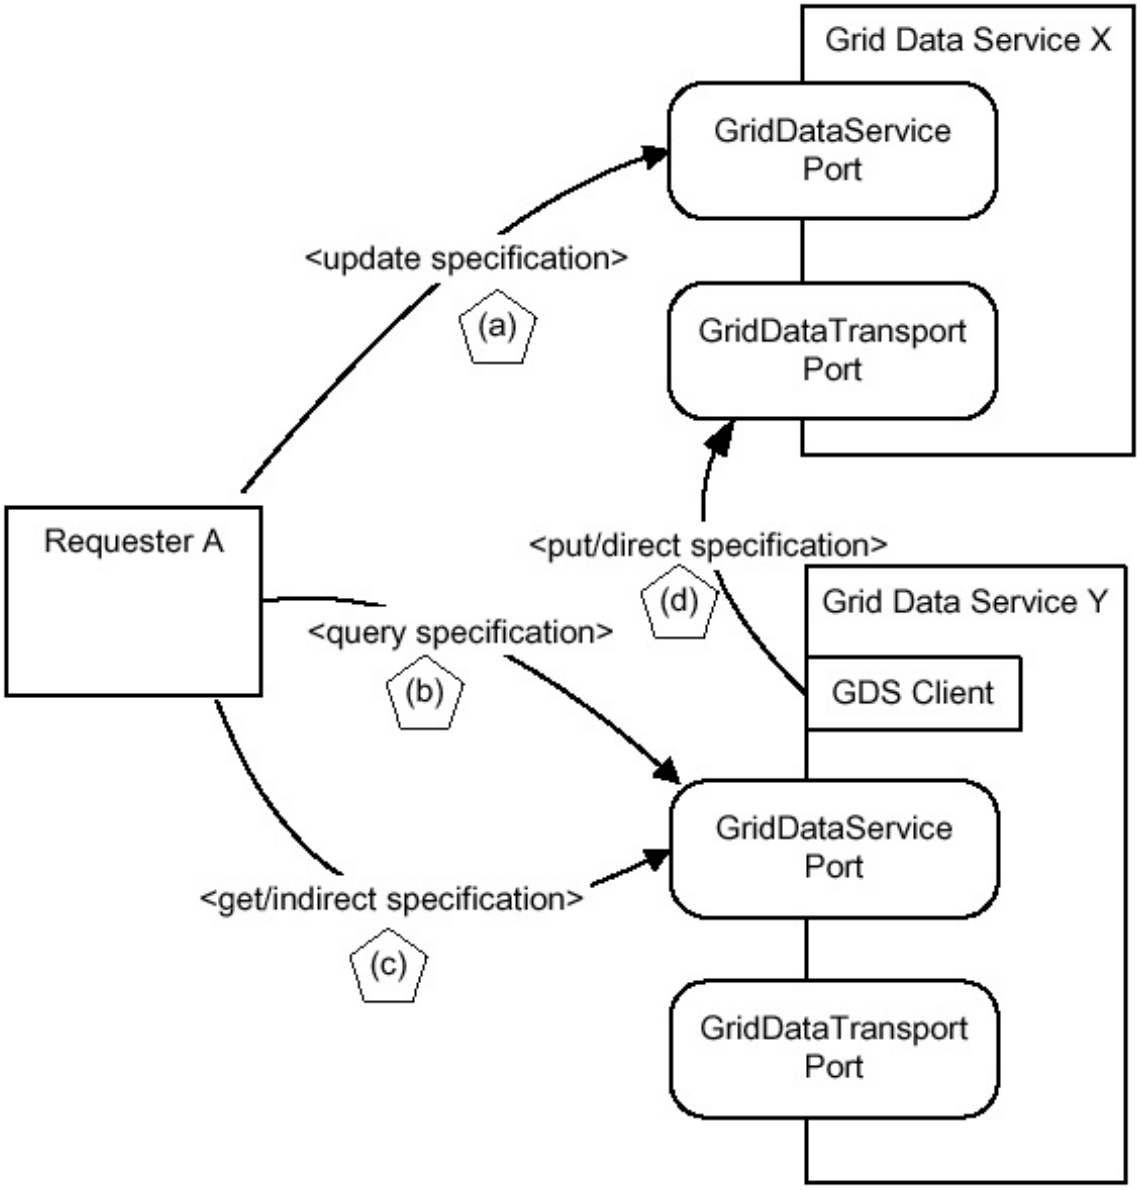

Source: Grid Database Service Specific

IBM Technology Directions

Federated DBMS Architecture

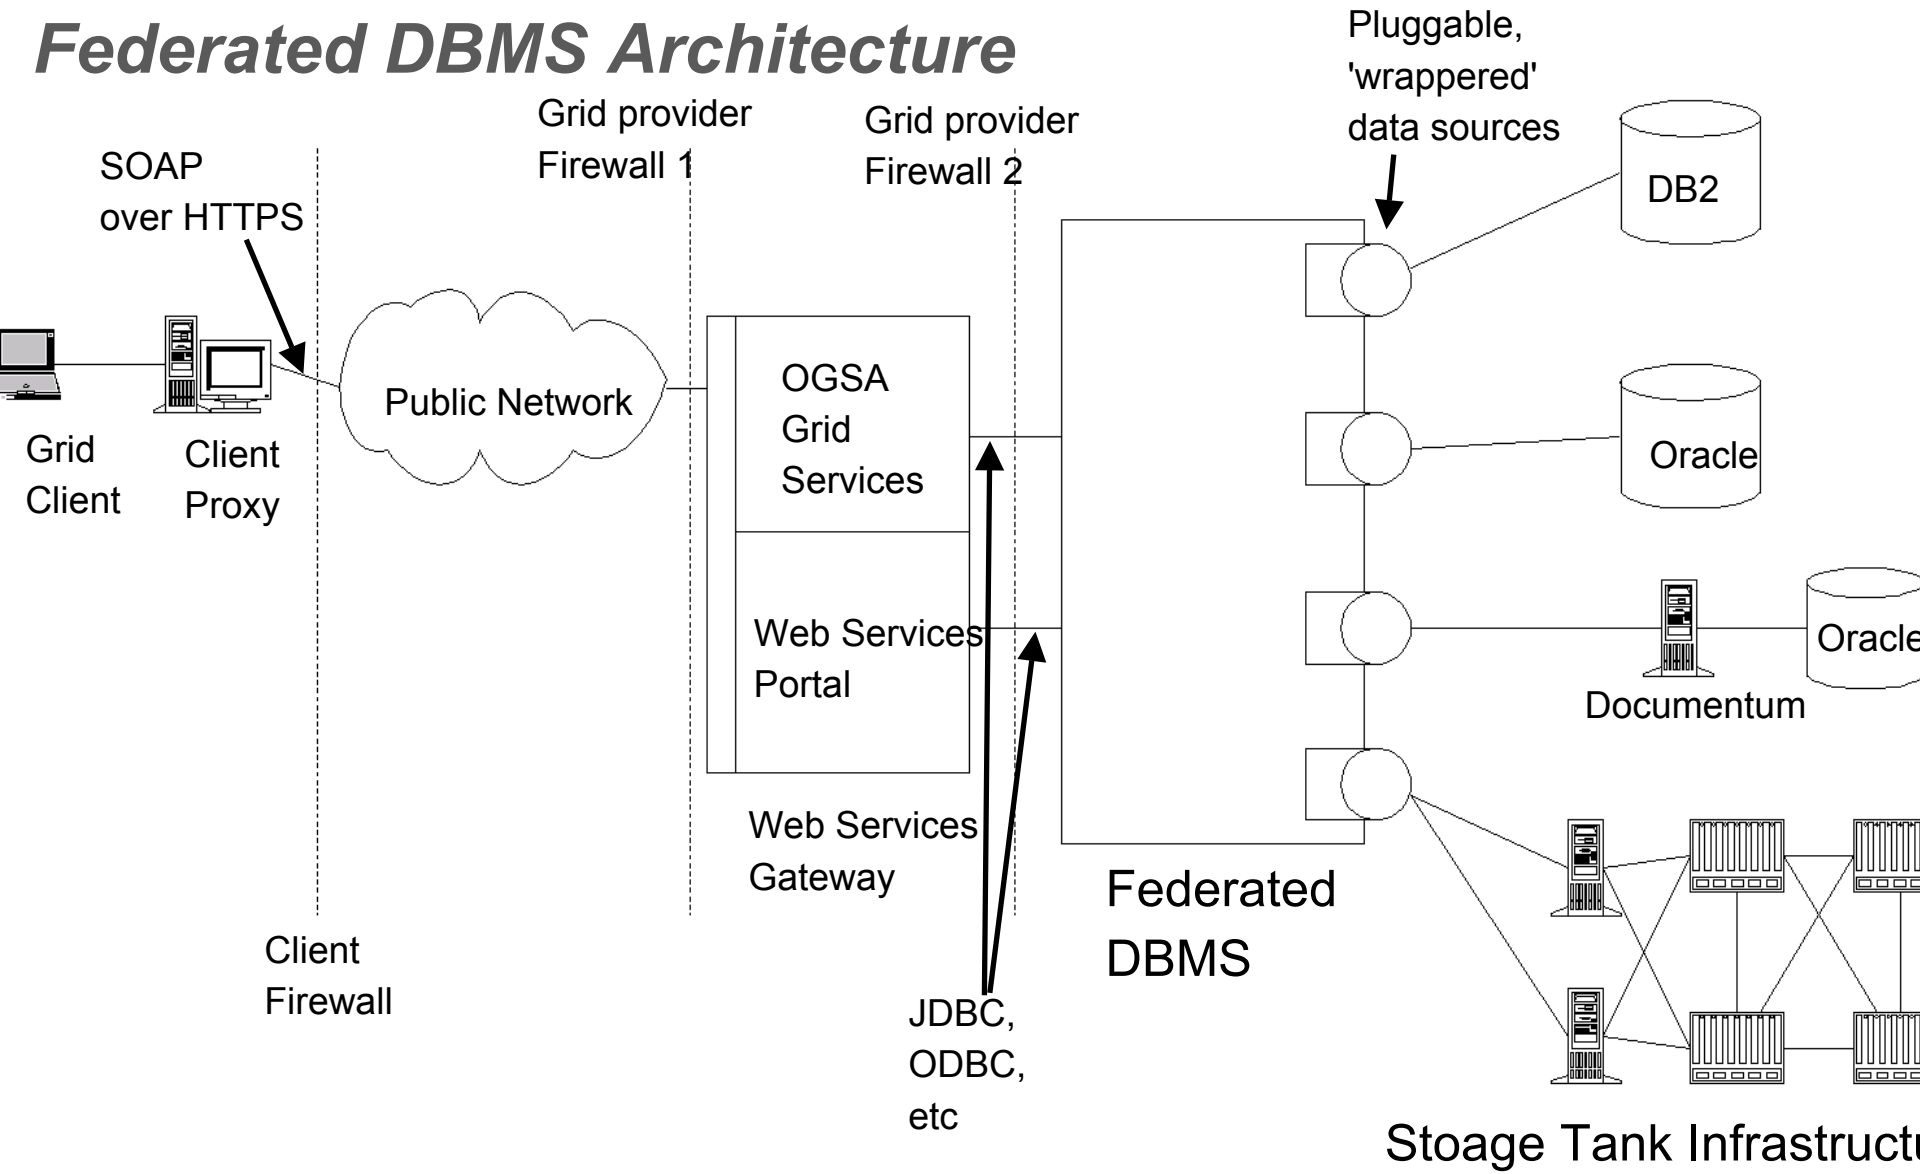

Globus Project and Toolkit

Globus Project

- **At its core, Globus is a research project. Globus research focuses not only on the issues associated with building computational grid infrastructures, but also on the problems that arise in designing and developing application that use grid services.**
- **Organized around four main activities.**
  - Research: study basic problems in areas such as resource management, security, information services, and data management.
  - Testbed: assist in planning and building large-scale testbeds, both for our own research and for production use by scientists and engineers.
  - Software Tools: We build robust research prototype software that runs on a variety of interesting and important platforms.
  - Applications: develop large-scale grid-enabled applications in collaboration with scientists and engineers.

ource: www.globus.org

Globus Toolkit™

- The Globus Project provides software tools that make it easier to build computational grids and grid-based applications. These tools are collectively called the Globus Toolkit™.
- Is an open architecture, open source software toolkit.
- Is used by many organizations to build computational grids that support their applications.

ource: www.globus.org

lobus Toolkit™ Version 2.2 Layered Grid Architecture

oordinating multiple resources”:
iquitous infrastructure services,
p-specific distributed services

haring single resources”:
gotiating access, controlling use

alking to things”:
mmunication (Internet
otocols) & security
ontrolling things locally”: Access
& control of, resources

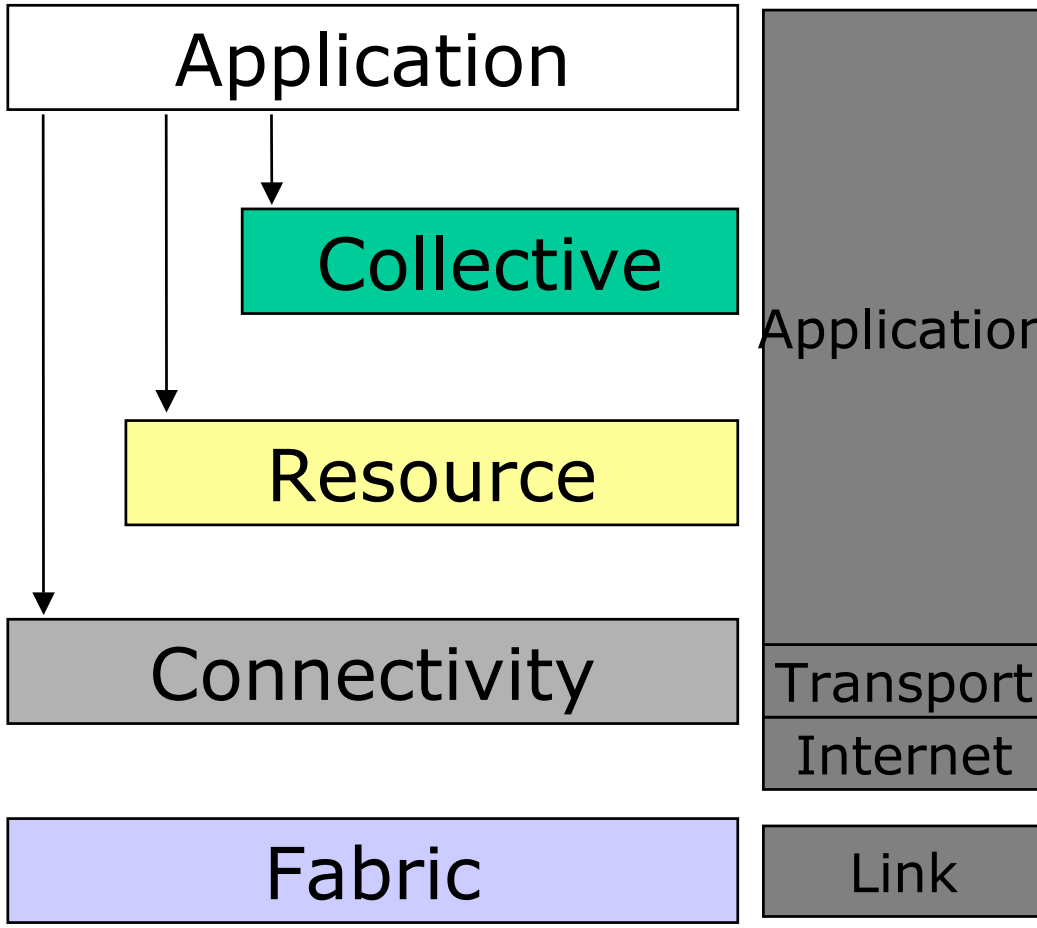

The Anatomy of the Grid: Enabling Scalable Virtual Organizations”, Foster, Kesselman, Tuecke, Intl Journal of High Performance Computing Applications, 15(3), 2001.

Globus Toolkit™ Version 2.2 Key Protocols

- **The Globus Toolkit™ v2 (GT2) centers around four key protocols**
  - Connectivity layer:
    - *Security*: Grid Security Infrastructure (GSI)
  - Resource layer:
    - *Resource Management*: Grid Resource Allocation Management (GRAM)
    - *Information*: Grid Resource Information Protocol (GRIP/LDAP)
    - *Data Transfer*: Grid File Transfer Protocol (GridFTP)
- **Also key collective layer protocols**
  - Monitoring & Discovery, Replication, etc.

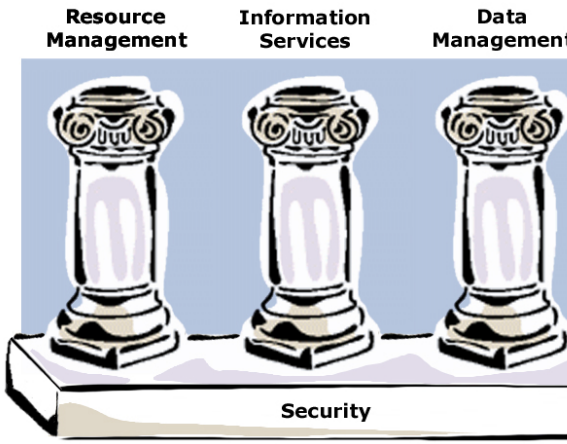

Globus Toolkit 2 Layered Grid Architecture

Protocols, Services, and APIs

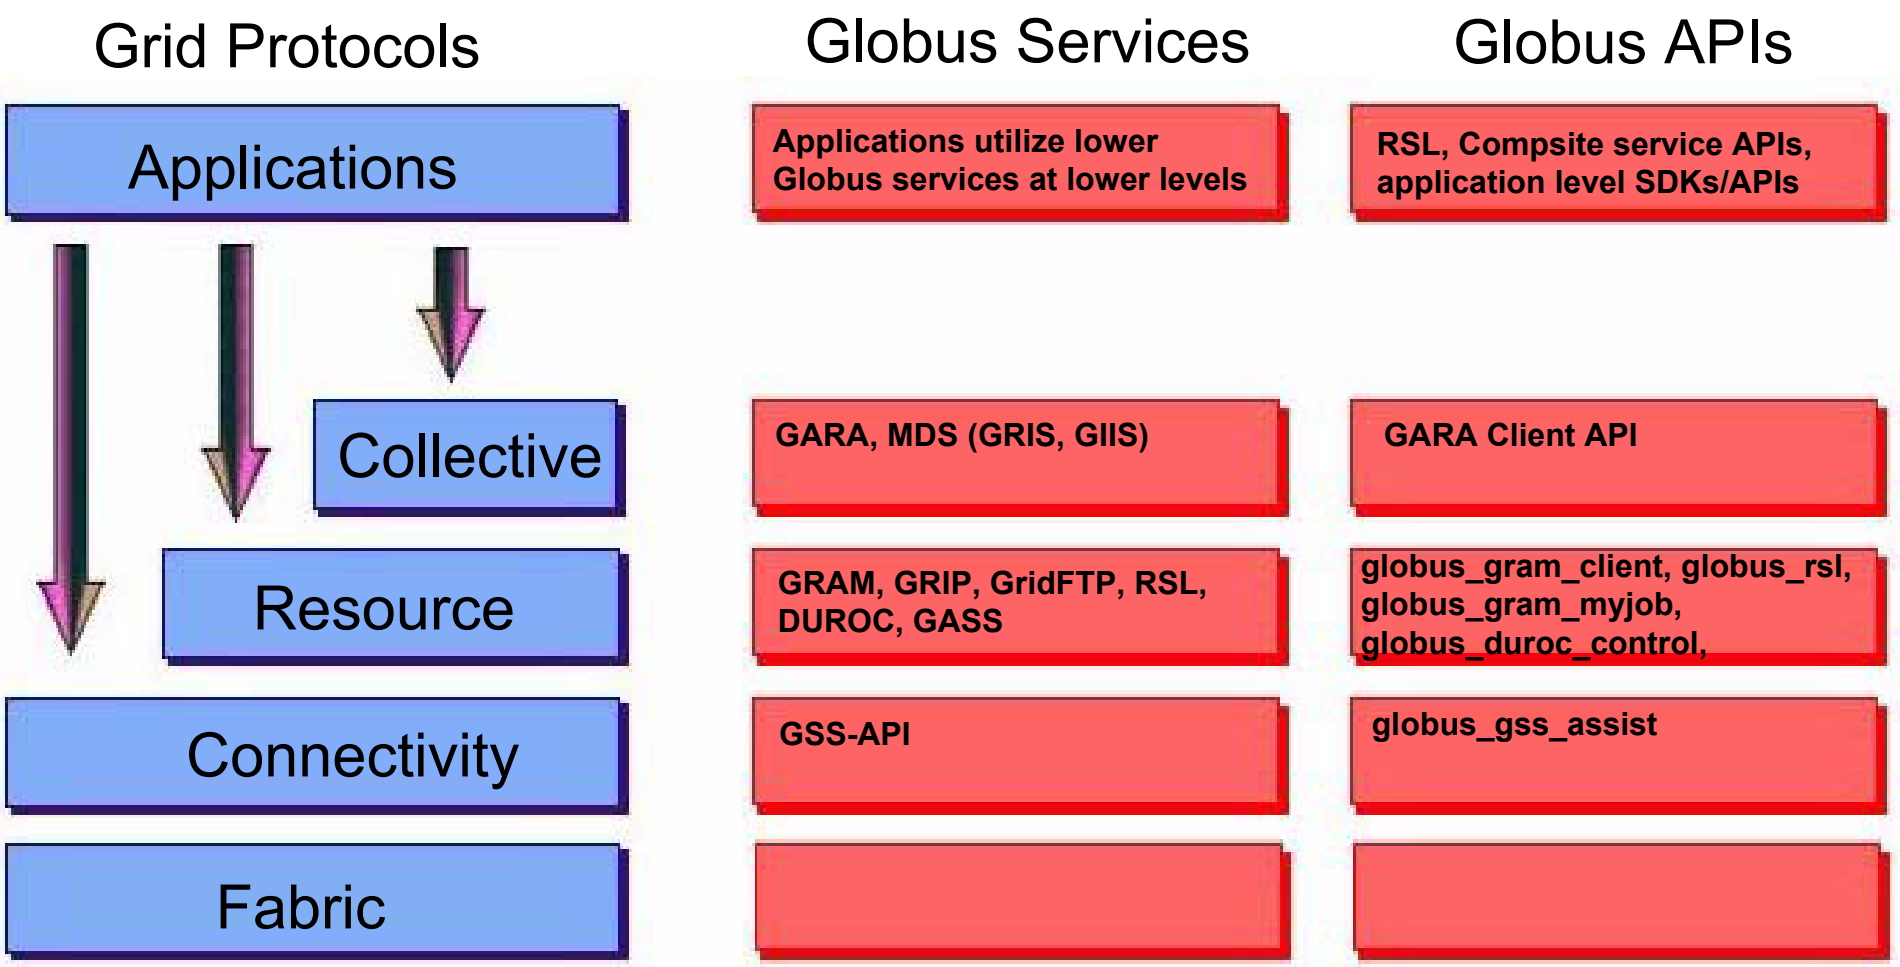

GT3 Architecture Overview

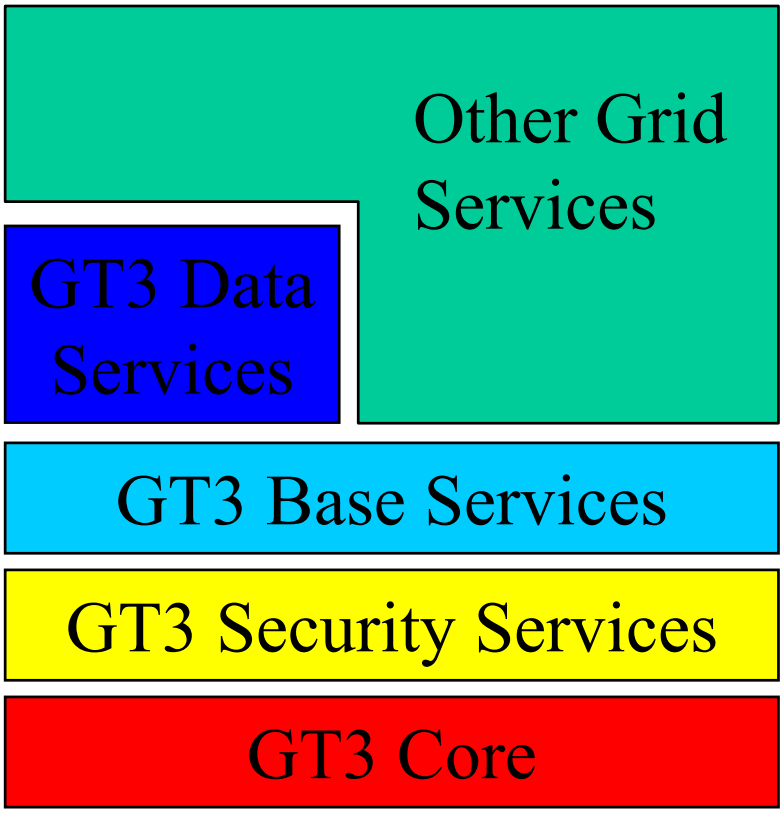

plica Management

anaged Job Service
Index Service

cure Conversation
rvice

idService

Workload Management

Diagnostics

File Streaming Service

Reliable File Transfer Service

NotificationSink

Registration

HandleResolver

NotificationSubscription

NotificationSource

Factory

Autonomic Computing

Autonomic Vision

- *"Intelligent" open systems that...*
  - Hide complexity
  - "Know" themselves
  - Adapt to unpredictable conditions
  - Continuously tune to meet performance goals
  - Recover from failures
  - Provide a safe environment
- Providing customers with...
  - Increased return on IT investment
  - Improved resiliency
  - Accelerated implementation of new capabilities

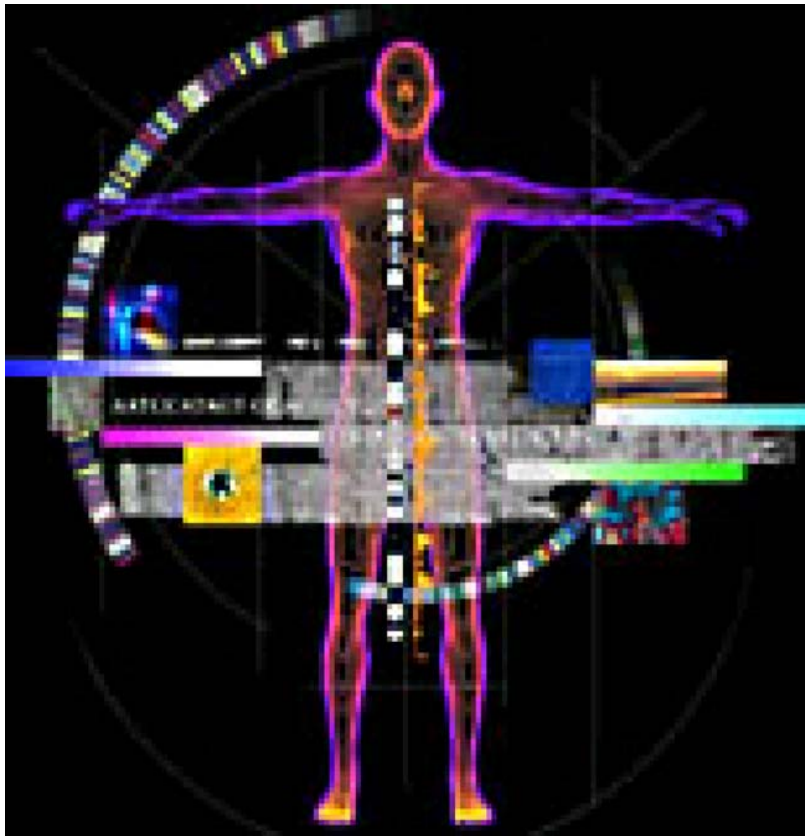

Autonomic Computing

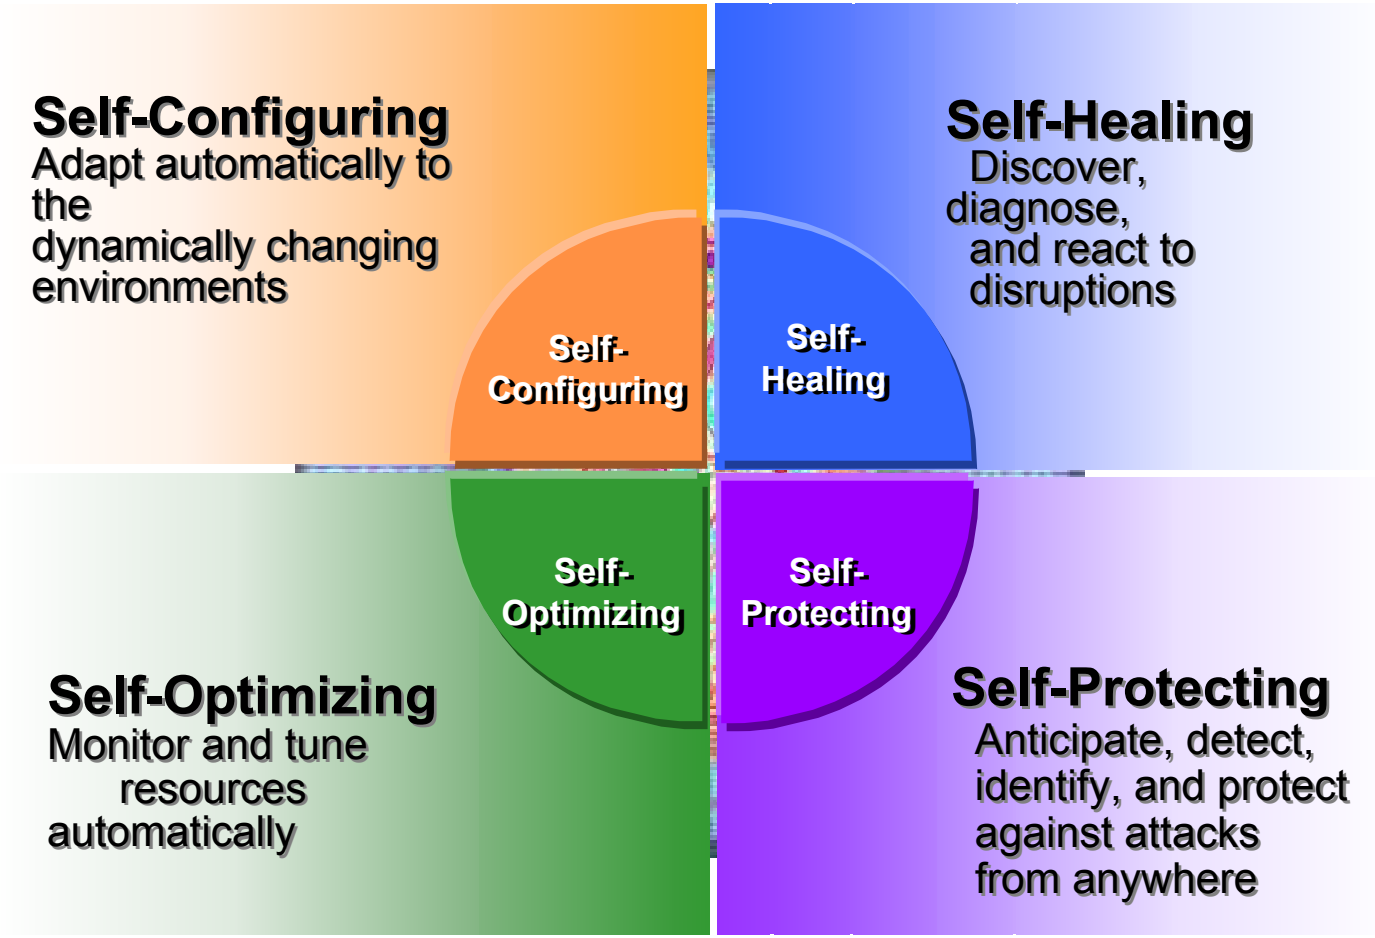

Autonomic Element

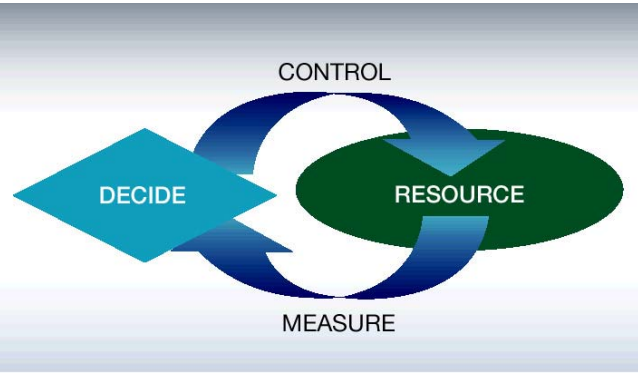

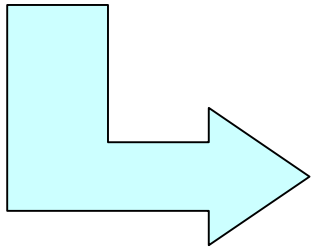

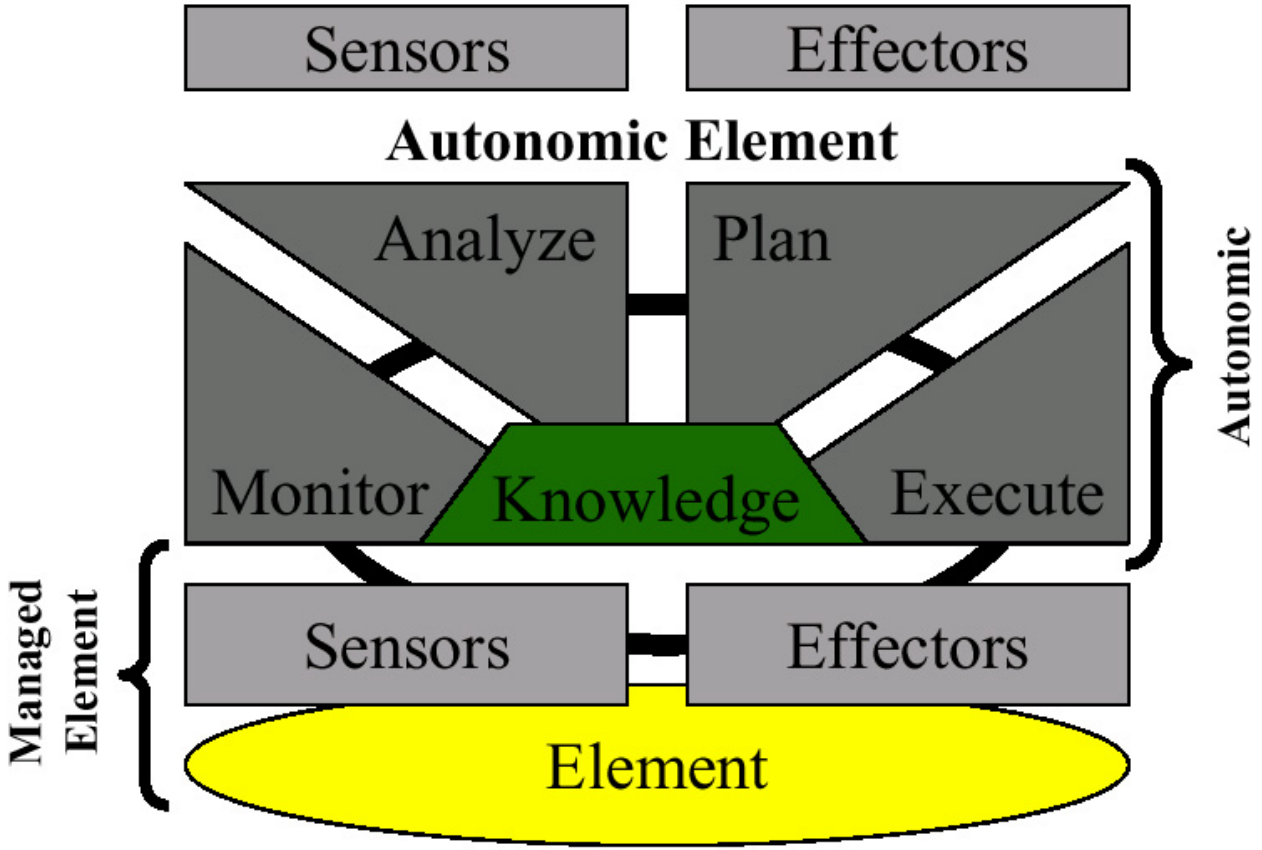

Autonomic Components in a Hierarchy

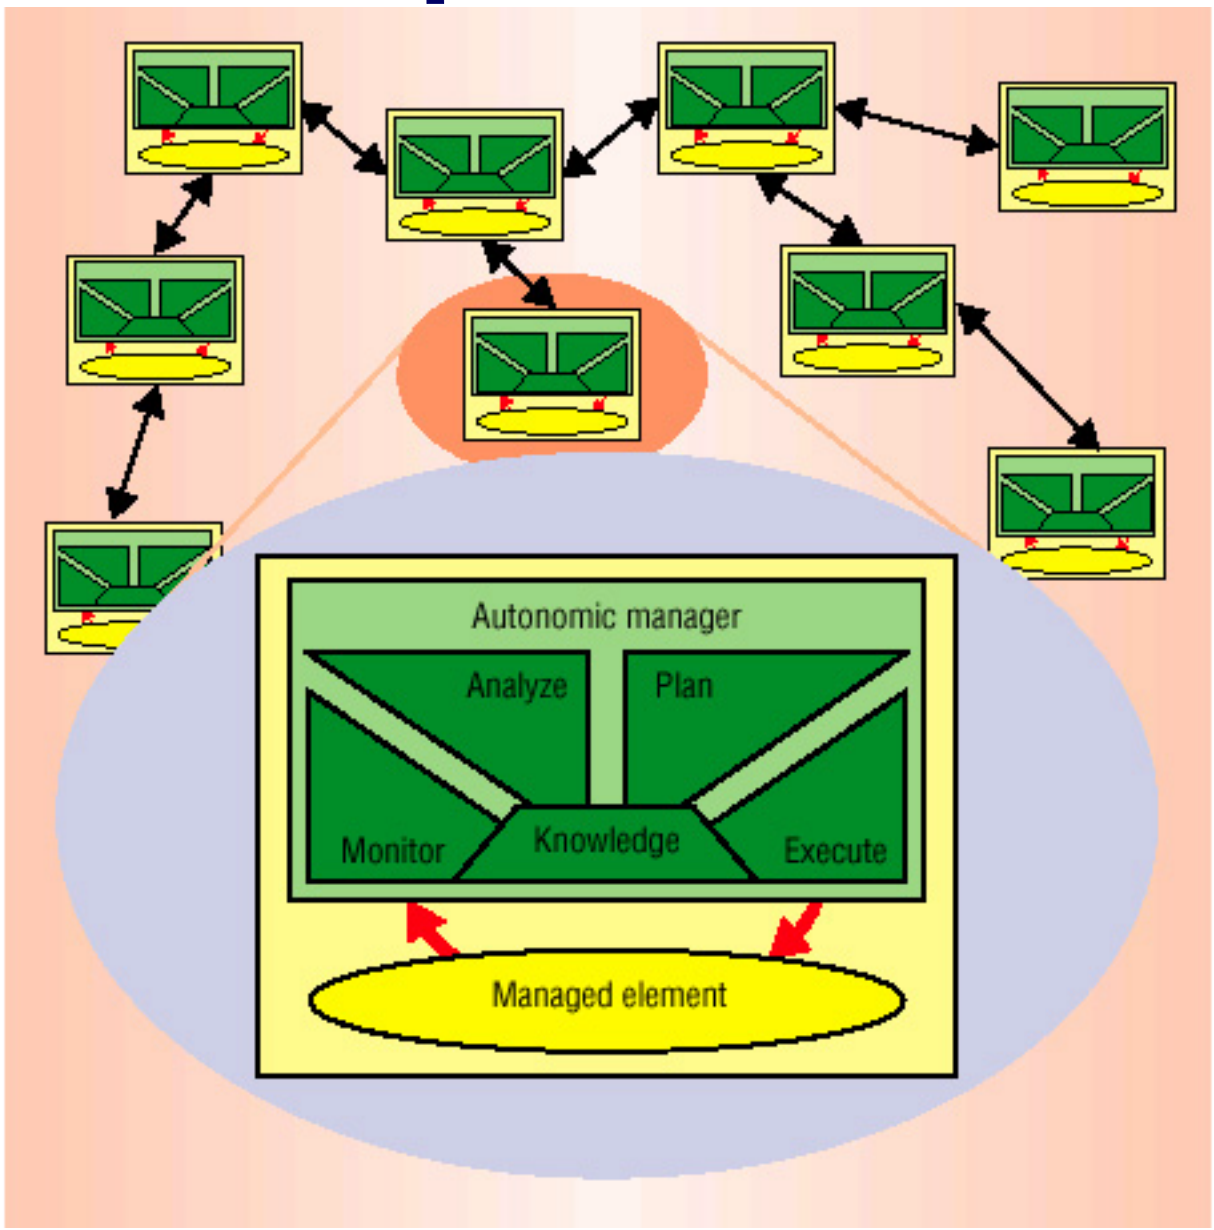

Self-Configuring Example: DB2 Configuration Advisor

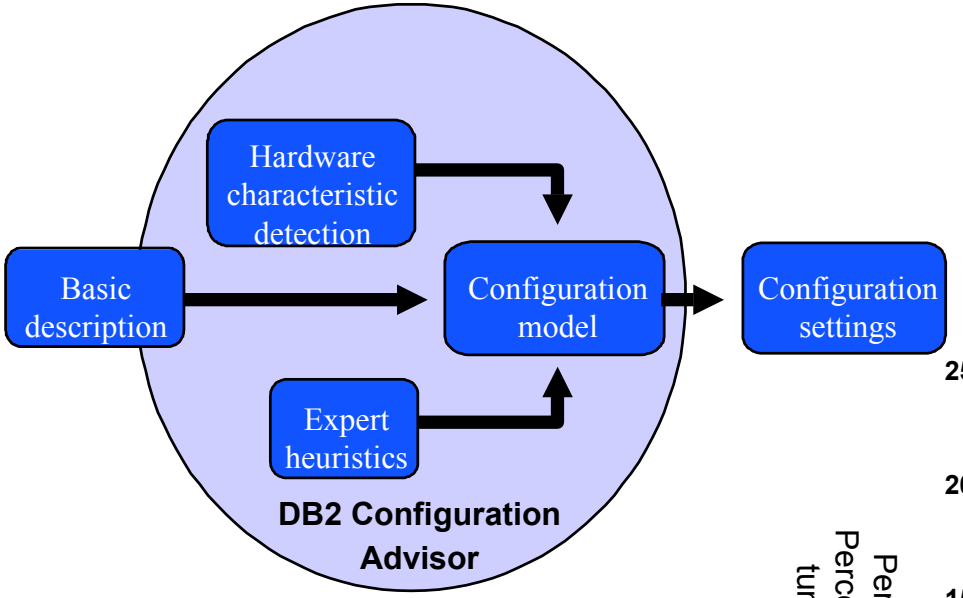

Speeds deployment

mproves performance

Frees up resource

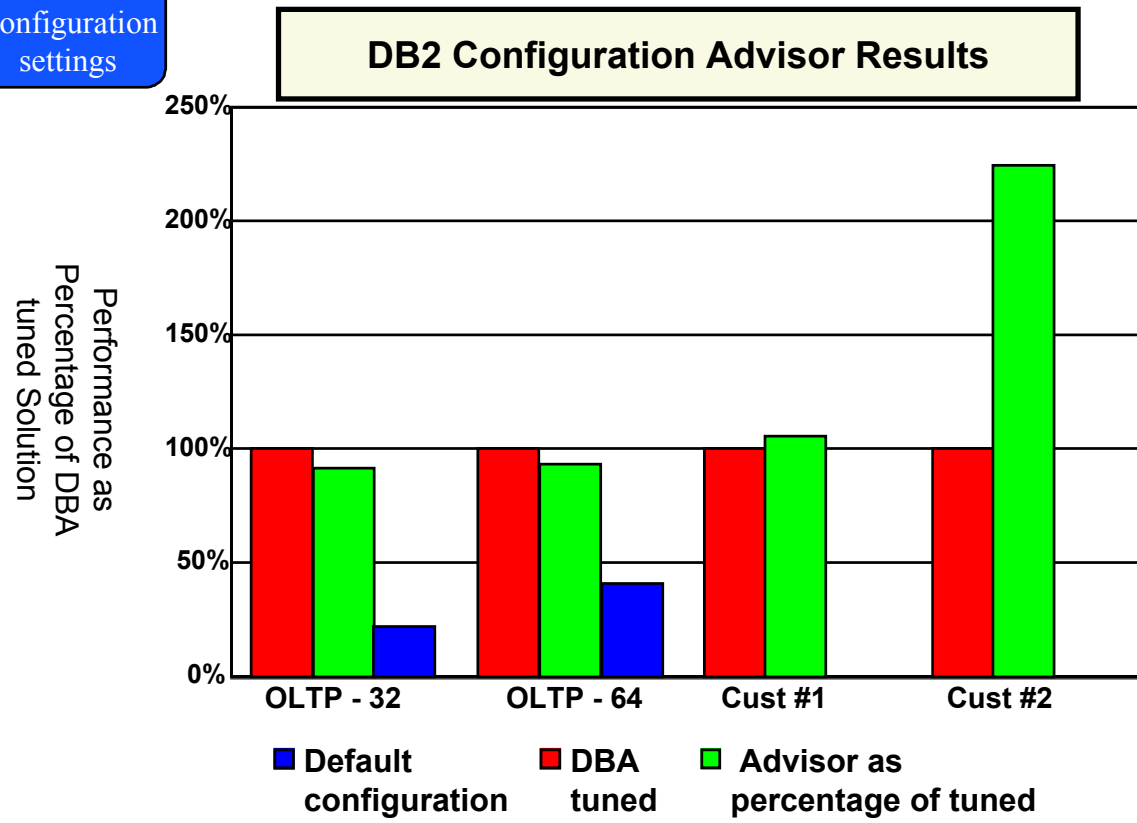

Autonomic Examples

stems nagement	 Tivoli. software	✓Access / Identity Managers ✓Storage Resource Manager ✓Service Level Advisor
ent	ThinkCentre ThinkPad ThinkVision ThinkVantage Technologies Think Accessories and Services ThinkVantage Design	✓ImageUltra ✓Rapid Restore PC ✓Embedded Security Subsystem
plication	 WebSphere. software	✓Prioritization of User Transactions ✓Custom Advisors ✓Problem Analysis and Recovery
tabase & llaboration	Product Store Month DB2. Data Management Software Lotus software	✓DB2 Query Patroller ✓Tivoli Analyzer for Domino
rvers	 eserver	✓Dynamic Partitioning ✓IBM Director ✓BladeCenter
orage	 TotalStorage™	✓Intelligent cache configuration ✓Predictive Failure Analysis ✓Dynamic volume expansion

Additional Information

Introduction to Grid Computing Video

- **Available at www.ibm.com/grid**
- **View online or download**

Content:

- What is Grid Computing
- Benefits of Grid Computing
- OGSA
- Customer Testimonials

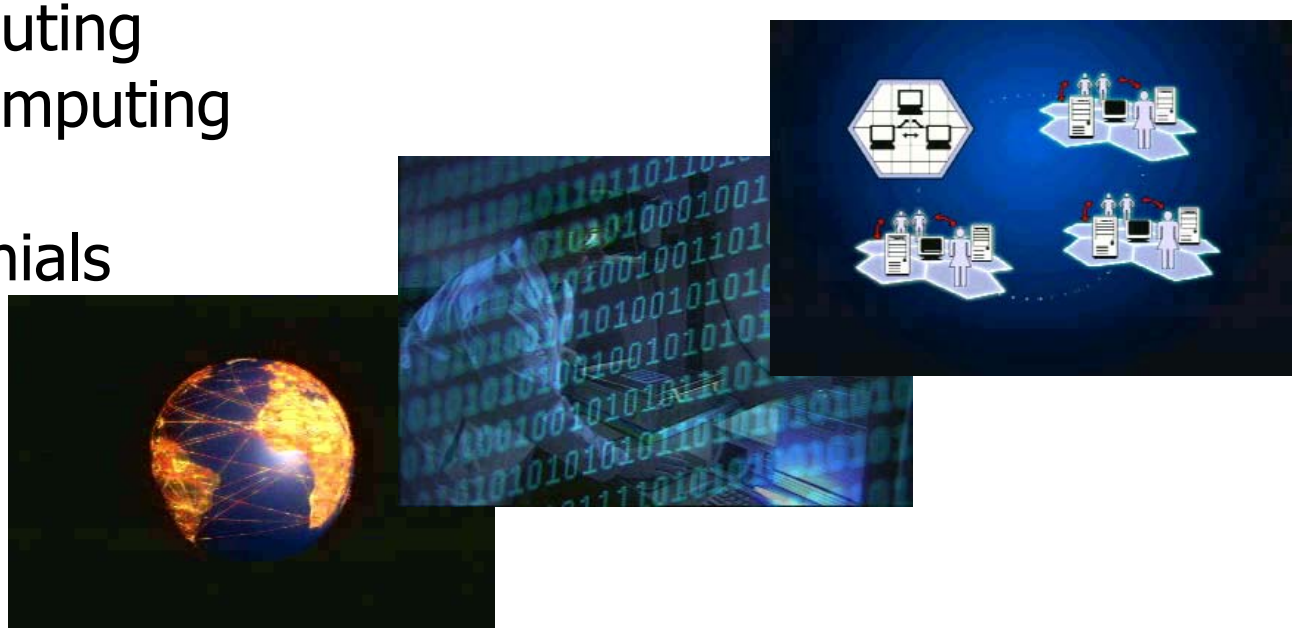

ITSO Redbook

- ***Redbook:*** **Introduction to Grid Computing with Globus**

- **Available**:
  - December 2002
  - Download from www.redbooks.ibm.com

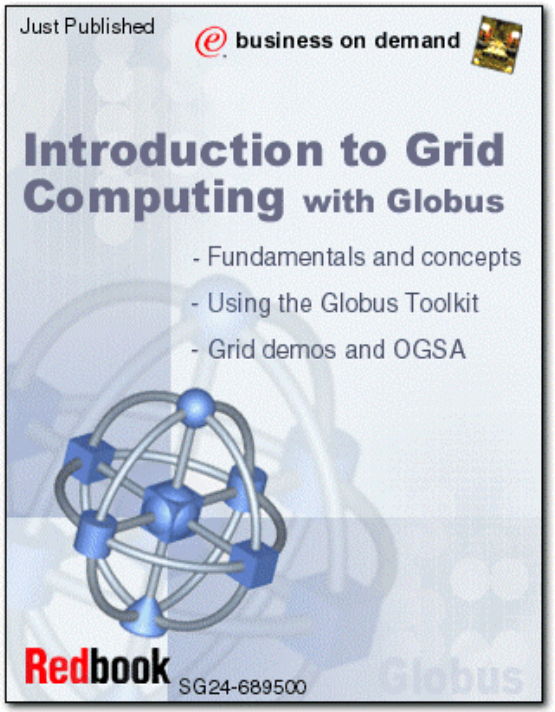

- Content:
  - Presents the architecture and components to design a Grid solution by using the Globus 2.0 Toolkit
  - Explains different Grid types
  - Architecture and security considerations
  - OGSA and Grid middleware
  - Showcases several real-life application examples

Learning Services Class

- ***Course:*** **Introduction to Grid Computing, the Globus Toolkit and OGSA**

- Content:
  - 2-day class, lecture-only
  - Based on the Globus tutorial of same name
  - Technical introduction both to Grid computing and the Globus Toolkit incl. descriptions of the core components
  - Usage of the Globus Toolkit in various applications
  - Future directions of Grid computing and the Globus Toolkit

More Courses planned for 2003 (e.g. Globus Developers+Admin Toolkits)

Grid and Autonomic Computing Information

- www.ibm.com/grid
- www.ibm.com/autonomic

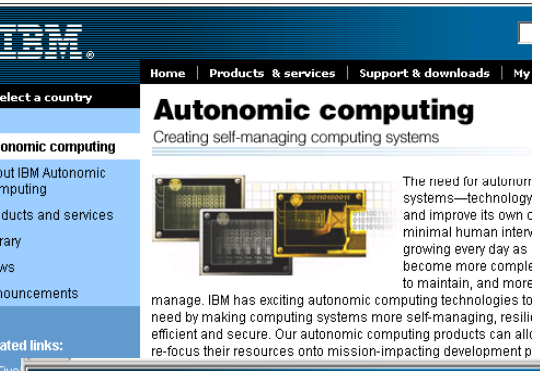

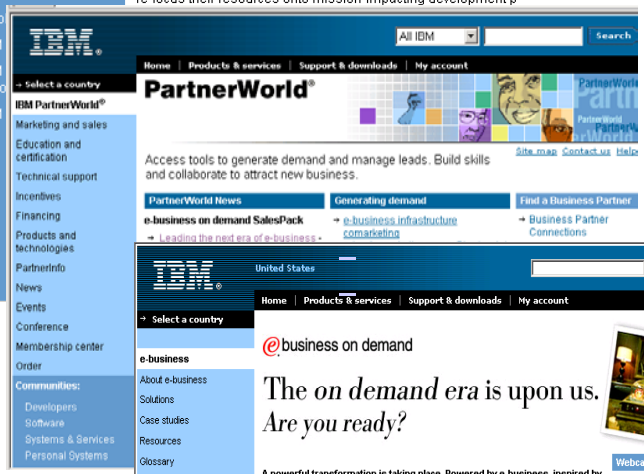

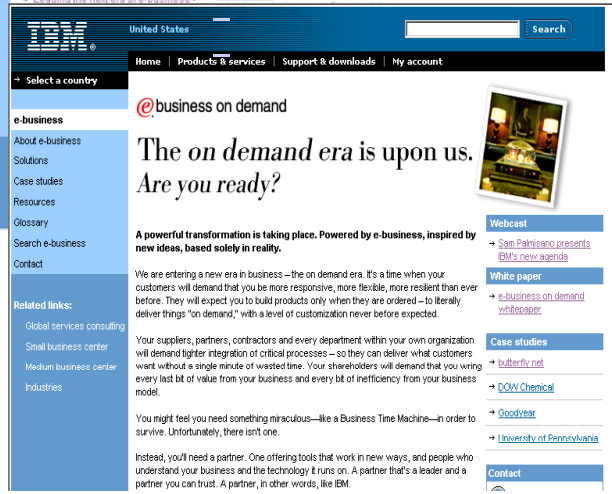

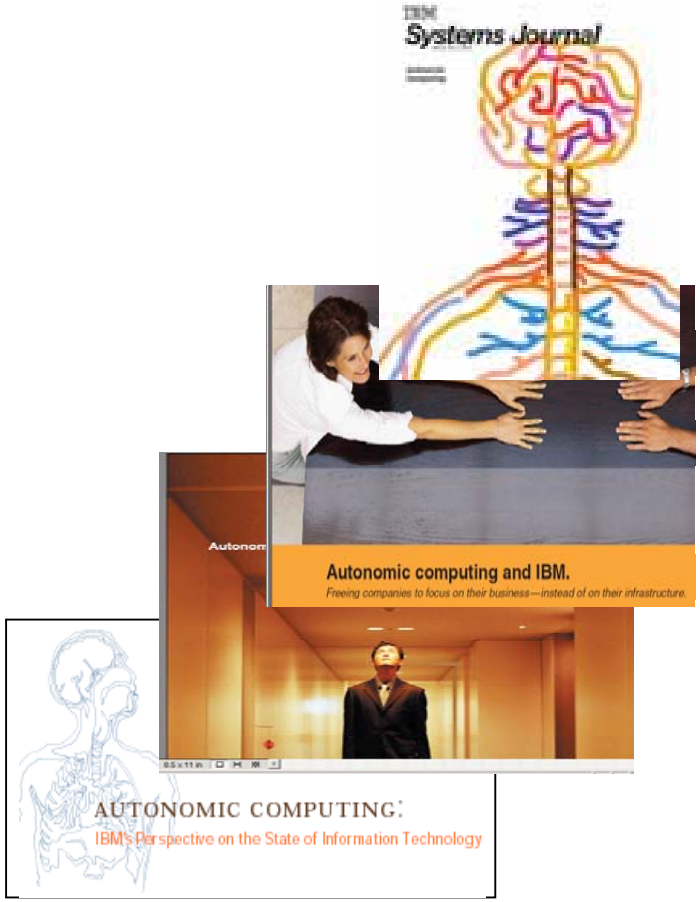

Questions?