The Grid
Enabling Resource Sharing within Virtual Organizations

Ian Foster

Mathematics and Computer Science Division
Argonne National Laboratory
and
Department of Computer Science
The University of Chicago

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

Invited Talk, WWW 2000 Conference, Honolulu, May 9, 2002

The Grid Vision

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

– On-demand, ubiquitous access to computing, data, and services
– New capabilities constructed dynamically and transparently from distributed services

"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)

Why the Grid? (1) Evolution of the Scientific Process

- Pre-electronic
  - Theorize &/or experiment, alone or in small teams; publish paper
- Post-electronic
  - Construct and mine very large databases of observational or simulation data
  - Develop computer simulations & analyses
  - Access specialized devices remotely
  - Exchange information quasi-instantaneously within distributed multidisciplinary teams

⇒ *Need to manage dynamic, distributed infrastructures, services, and applications*

eScience Application: Sloan Digital Sky Survey Analysis

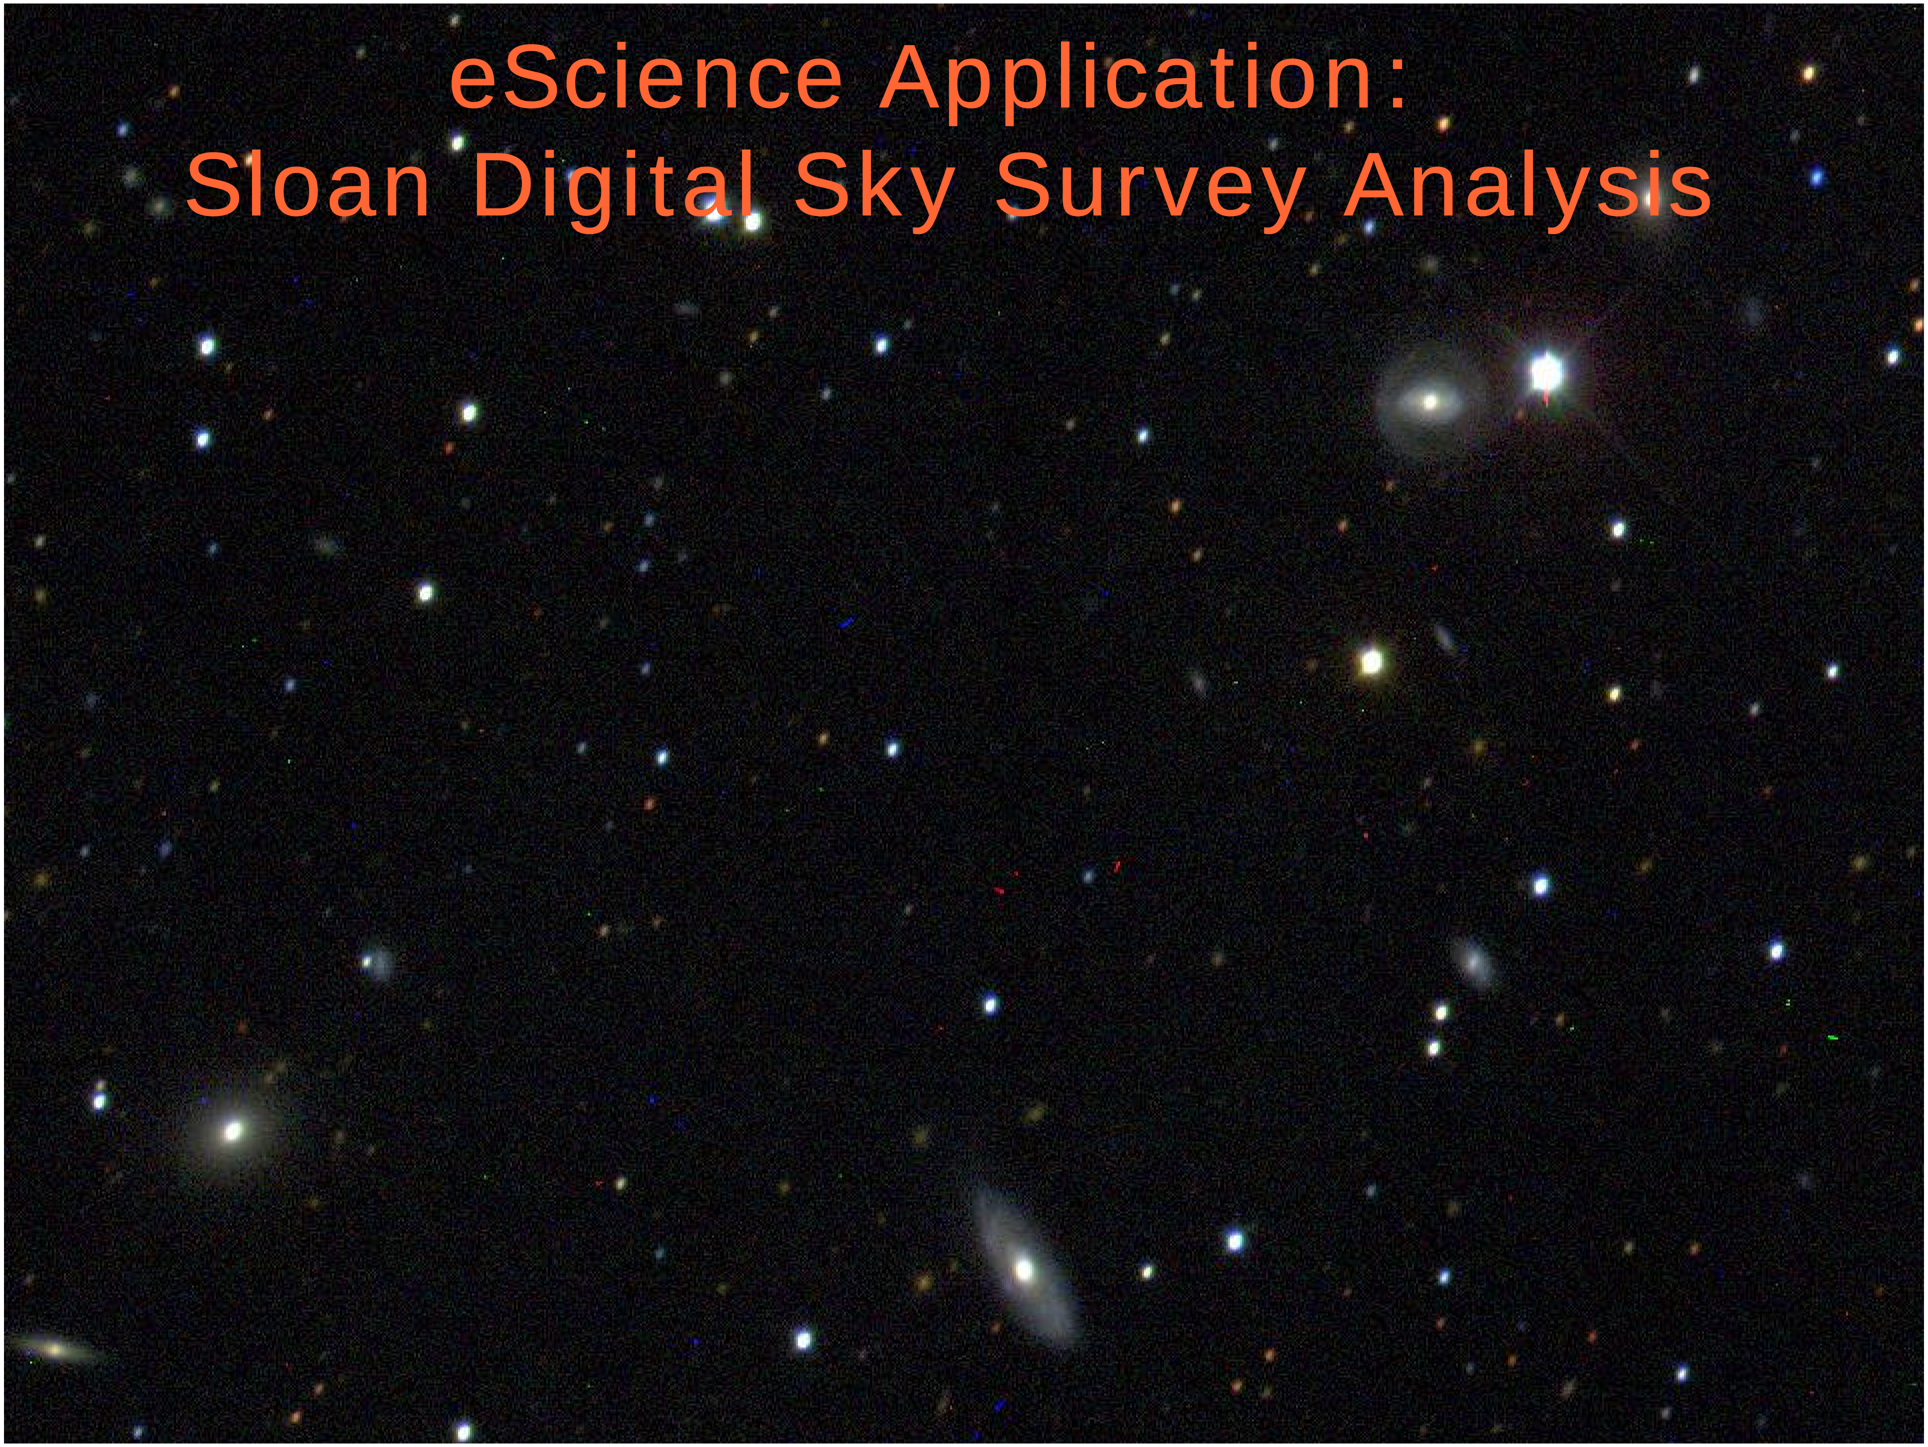

eScience Application: Sloan Digital Sky Survey Analysis

Size distribution of galaxy clusters?

eScience Application: Sloan Digital Sky Survey Analysis

Size distribution of galaxy clusters?

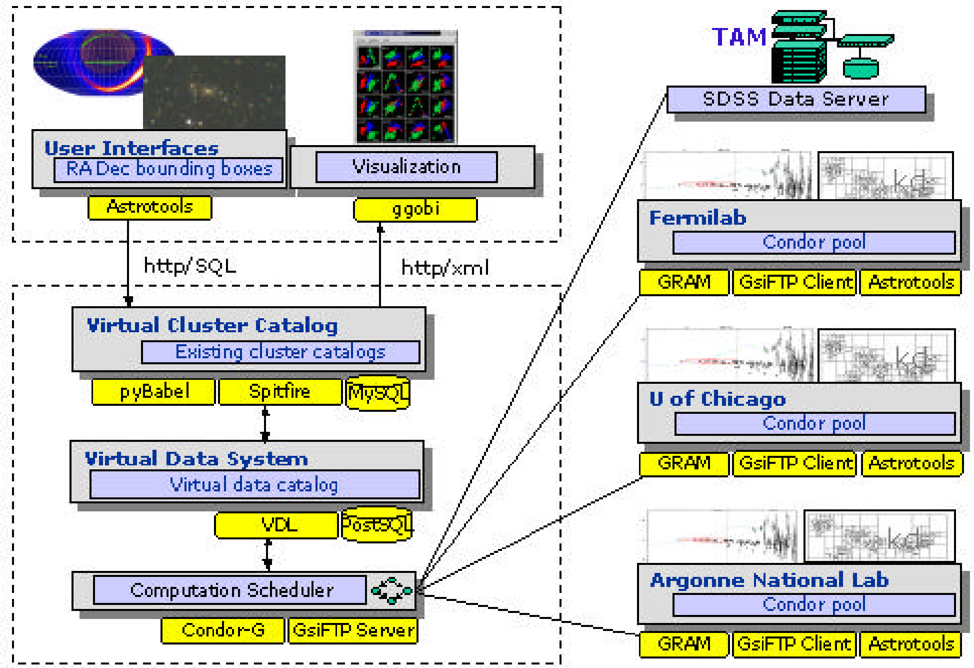

Chimera Virtual Data System
+ iVDGL Data Grid (many CPUs)

eScience Application: Sloan Digital Sky Survey Analysis

Size distribution of galaxy clusters?

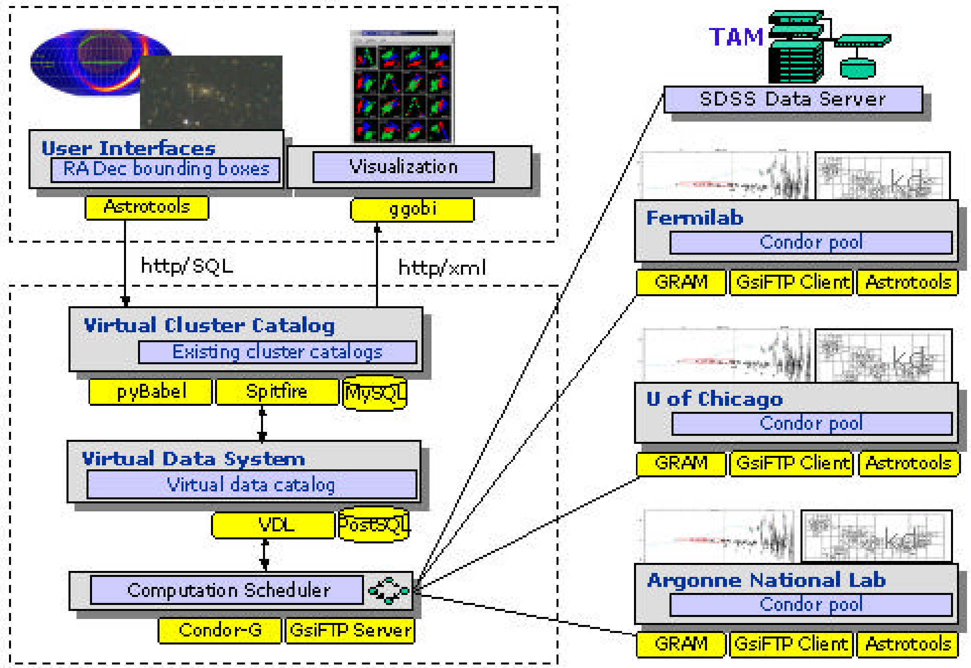

Chimera Virtual Data System + iVDGL Data Grid (many CPUs)

Galaxy cluster size distribution

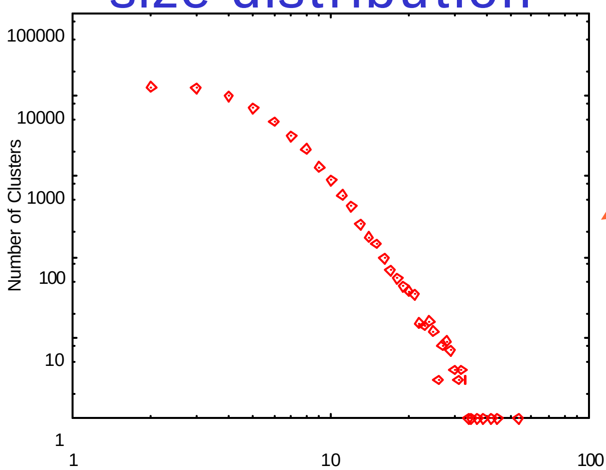

Why the Grid?
(2) Evolution of Business

- Pre-Internet
  - Central corporate data processing facility
- Post-Internet
  - Enterprise computing is highly distributed, heterogeneous, inter-enterprise (B2B)
  - Business processes computing- & data-rich
  - Outsourcing becomes feasible => service providers of various sorts

⇒ *Need to manage dynamic, distributed infrastructures, services, and applications*

Today's Enterprise Computing Environment

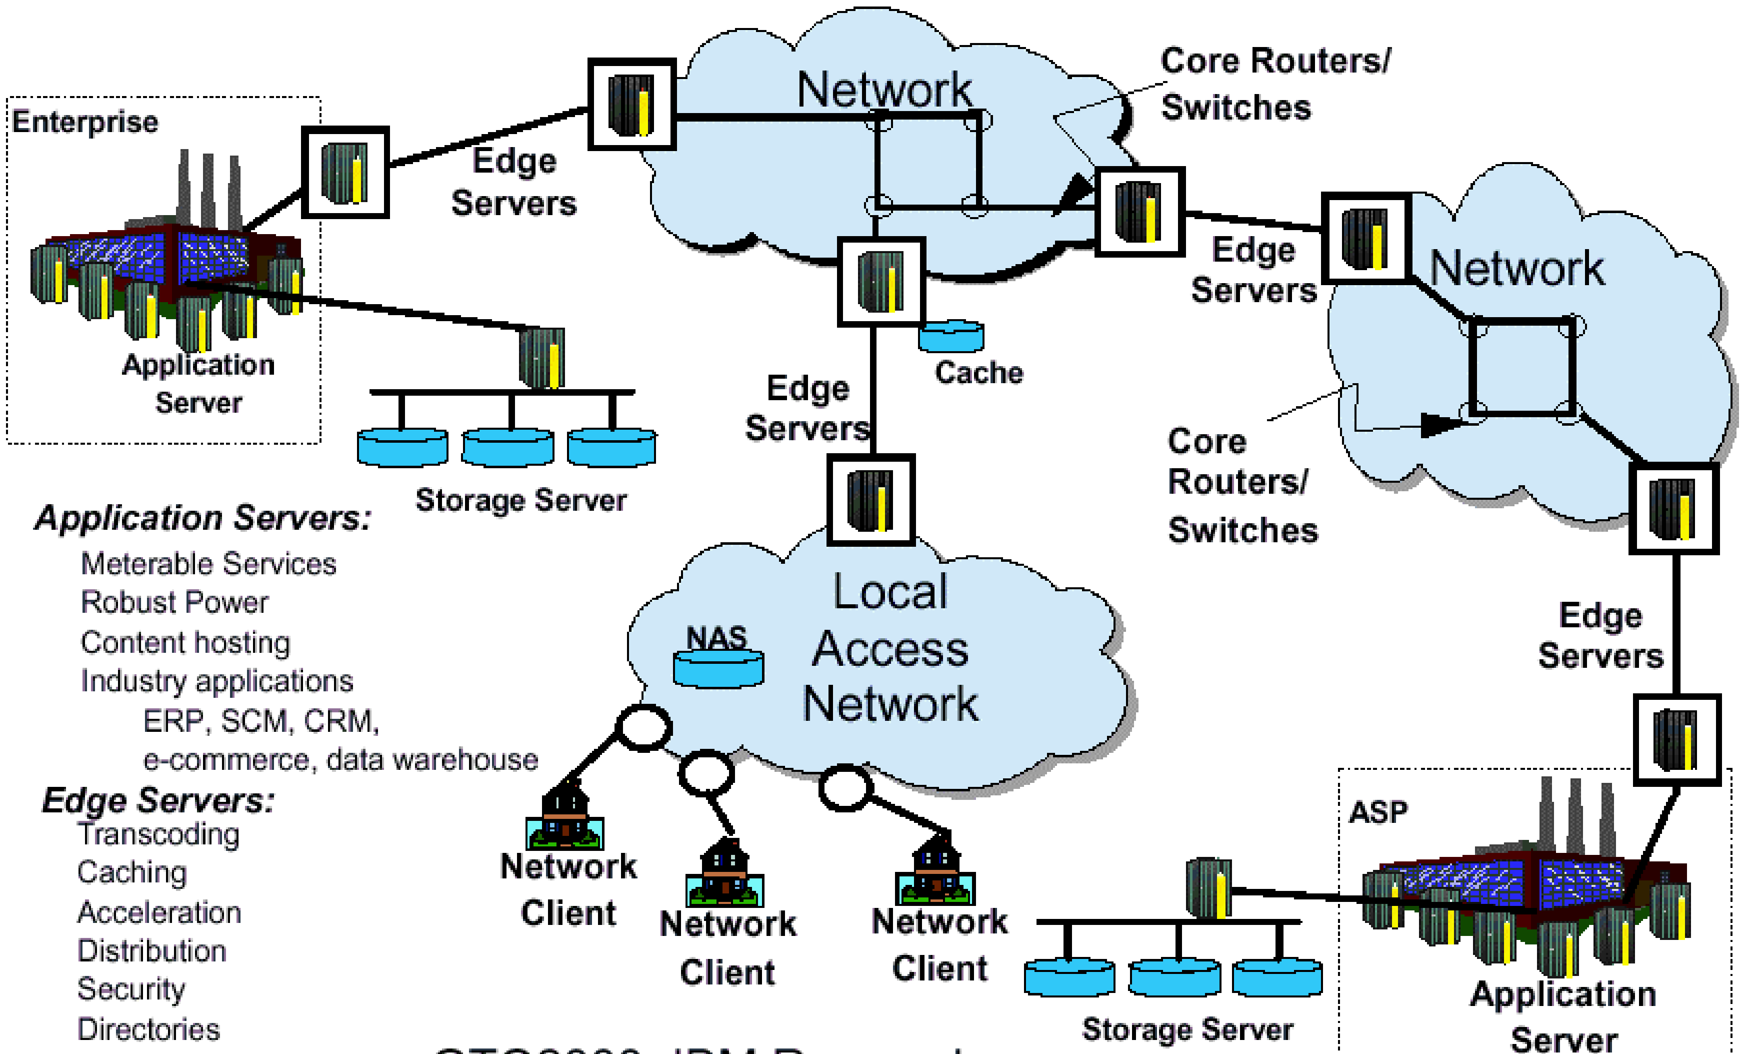

Application Servers:
- Meterable Services
- Robust Power
- Content hosting
- Industry applications
  - ERP, SCM, CRM,
  - e-commerce, data warehouse

Edge Servers:
- Transcoding
- Caching
- Acceleration
- Distribution
- Security
- Directories
- Quality of Service

GTO2000: IBM Research

Today's Enterprise Computing Environment

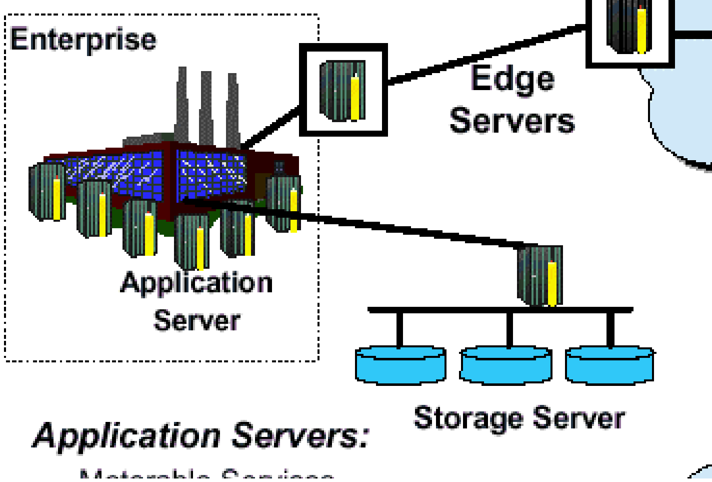
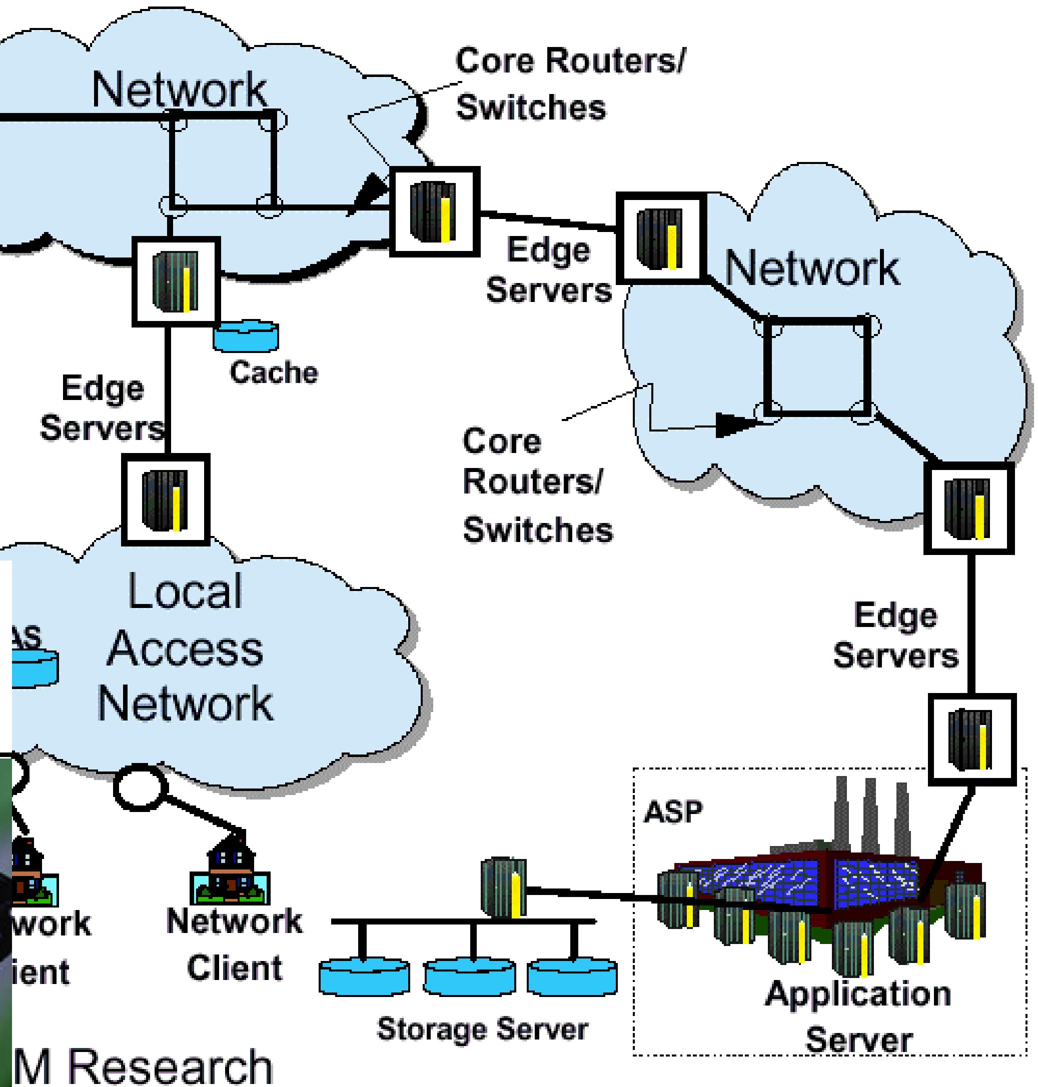

April 9th: Grid for multi-player games

Grid Computing

RED HERRING

TOP TEN TRENDS 2001

DISTRIBUTE THE WEALTH

Distributed computing initiatives.

GLOBUS PROJECT

Globus is a research and development project focused on enabling the application of grid computing concepts to scientific and engineering computing. The grid is an emerging infrastructure protocol that enables the integrated use of remote high-end computers, databases, scientific instruments, networks, and other resources.

Computing power on tap

Jun 21st 2001

From The Economist print edition

Economist.com

In the first of two articles, we look at the most ambitious attempt yet to combine millions of computers seamlessly around the world—to make processing power available on demand anywhere, rather like electrical power.

I.B.M. Making a Commitment to Next Phase of the Internet

The New York Times

By STEVE LOHR

I.B.M. is announcing today a new initiative to support and exploit a technology known as grid computing, which the company and much of the computer research community say is the next evolutionary step in the development of the Internet.

Merrill Lynch

Globus Grid Computing—the Next Internet

by John Roy/Steve Milunovich

The Internet was first a network and is now a communications platform. The next evolutionary step could be to a platform for distributed computing. This ability to manage applications and share data over the network is called "grid computing."

Challenging Technical Requirements

- Dynamic formation and management of virtual organizations
- Online negotiation of access to services: who, what, why, when, how
- Configuration of applications and systems able to deliver multiple qualities of service
- Autonomic management of distributed infrastructures, services, and applications
- Management of distributed state as a fundamental issue

State of the Art: Globus Toolkit™ (since 1996)

- Small, standards-based set of protocols for distributed system management
  - Authentication, delegation; resource discovery; reliable invocation; etc.
- Information-centric design
  - Data models; publication, discovery protocols
- Open source implementation
  - Large international user community
- Successful enabler of higher-level services and applications

Grid Projects in eScience: Rich in Ideas, Impact, and Logos

Grid Evolution: Open Grid Services Architecture

- Refactor Globus protocol suite to enable common base and expose key capabilities
- Service orientation to virtualize resources and unify resources/services/information
- Embrace key Web services technologies for standard IDL, leverage commercial efforts
- Result: standard interfaces & behaviors for distributed system management: the Grid service

Open Grid Services Architecture: Transient Service Instances

- "Web services" address discovery & invocation of persistent services
  - Interface to persistent state of entire enterprise
- In Grids, must also support transient service instances, created/destroyed dynamically
  - Interfaces to the states of distributed activities
  - E.g. workflow, video conf., dist. data analysis
- Significant implications for how services are managed, named, discovered, and used
  - In fact, much of OGSA (and Grid) is concerned with the management of service instances

Open Grid Services Architecture

- Defines fundamental (WSDL) interfaces and behaviors that define a Grid Service
  - Required + optional interfaces = WS “profile”
  - A unifying framework for interoperability & establishment of total system properties
- Defines WSDL extensibility elements
  - E.g., serviceType (a group of portTypes)
- Delivery via open source Globus Toolkit 3.0
  - Leverage GT experience, code, community
- And commercial implementations

The Grid Service = Interfaces/Behaviors + Service Data

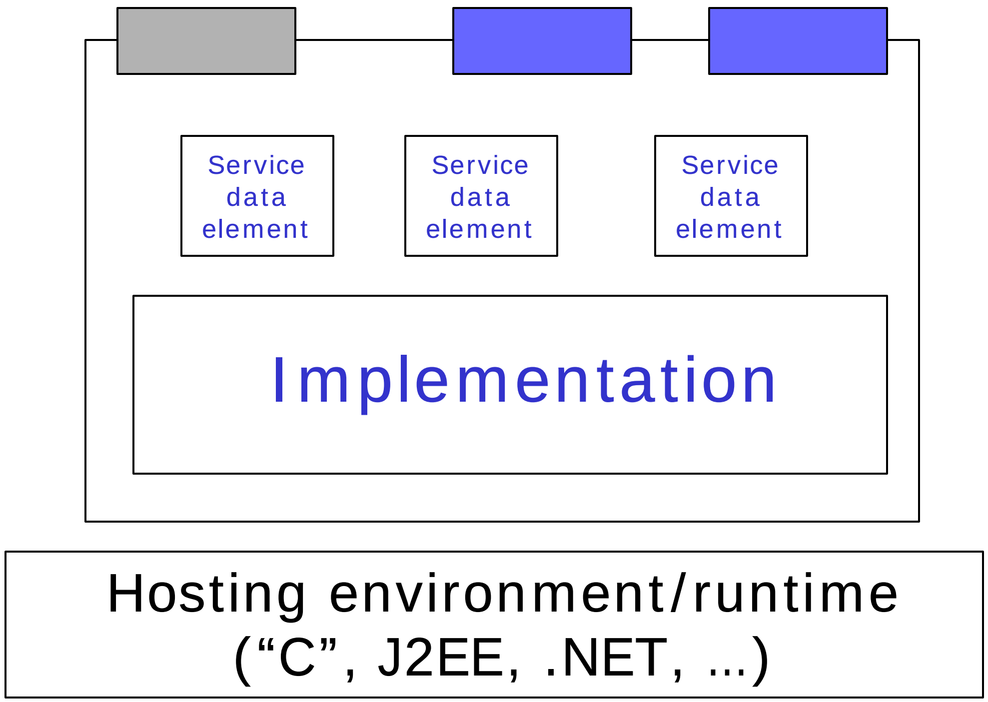

The Grid Service = Interfaces/Behaviors + Service Data

Service data access
Explicit destruction
Soft-state lifetime

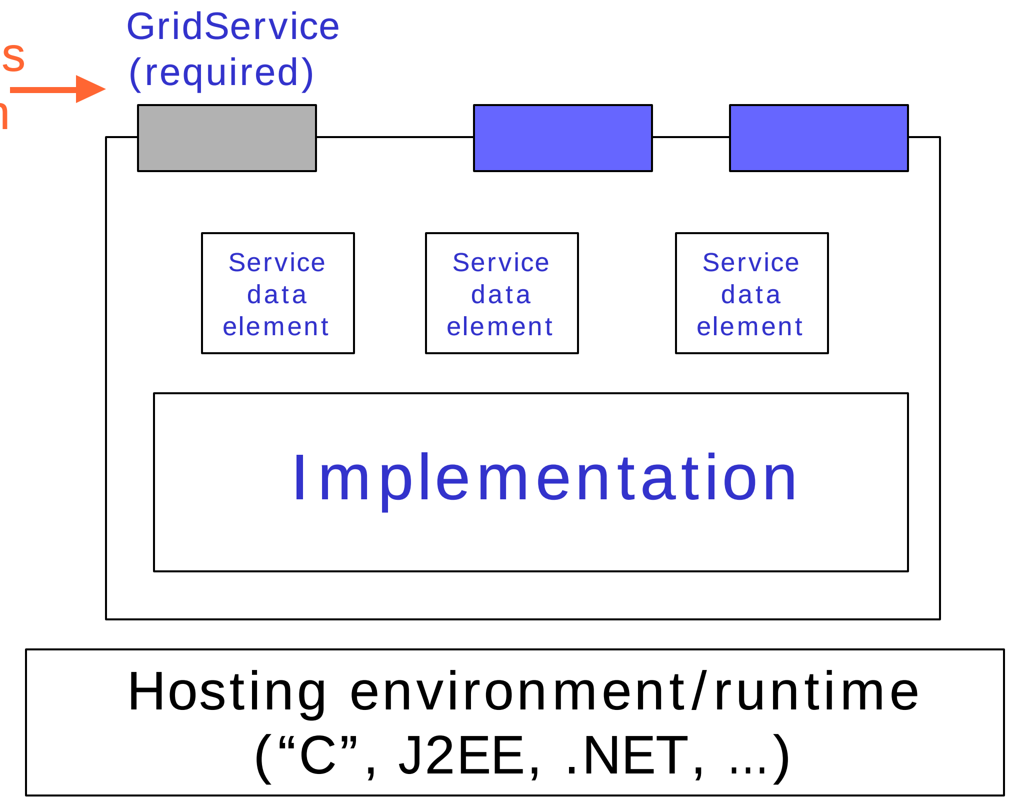

The Grid Service = Interfaces/Behaviors + Service Data

Service data access
Explicit destruction
Soft-state lifetime

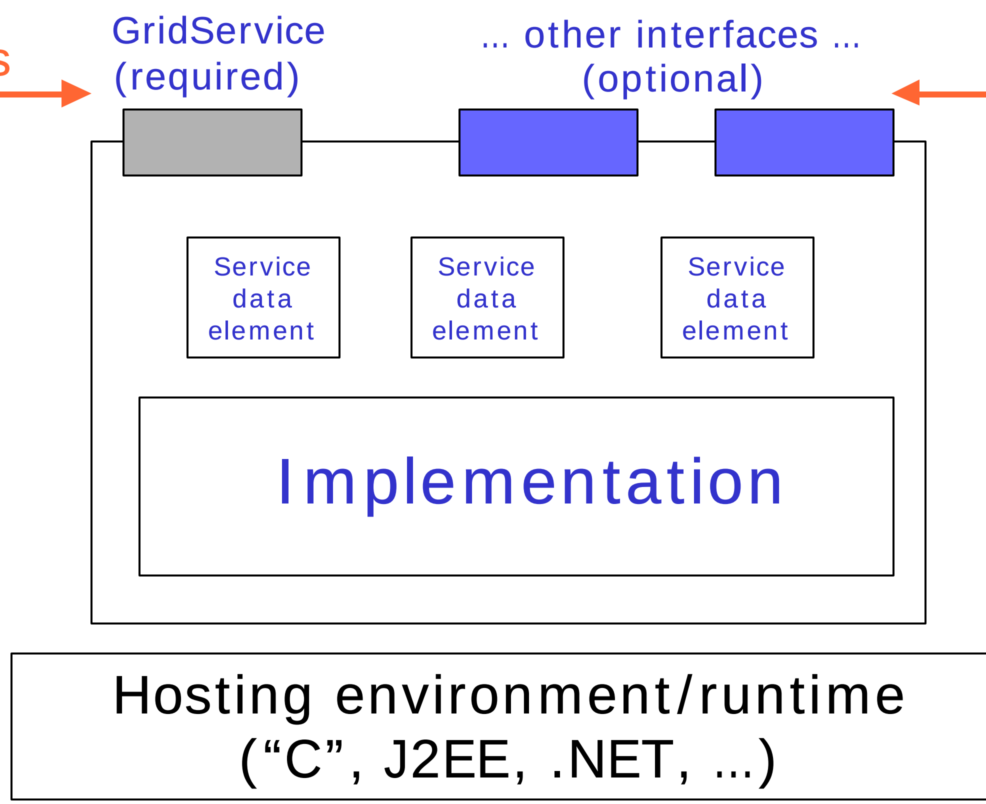

Standard:
- Notification
- Authorization
- Service creation
- Service registry
- Manageability
- Concurrency

+ application-specific interfaces

The Grid Service = Interfaces/Behaviors + Service Data

Service data access
Explicit destruction
Soft-state lifetime

Binding properties:
- Reliable invocation
- Authentication

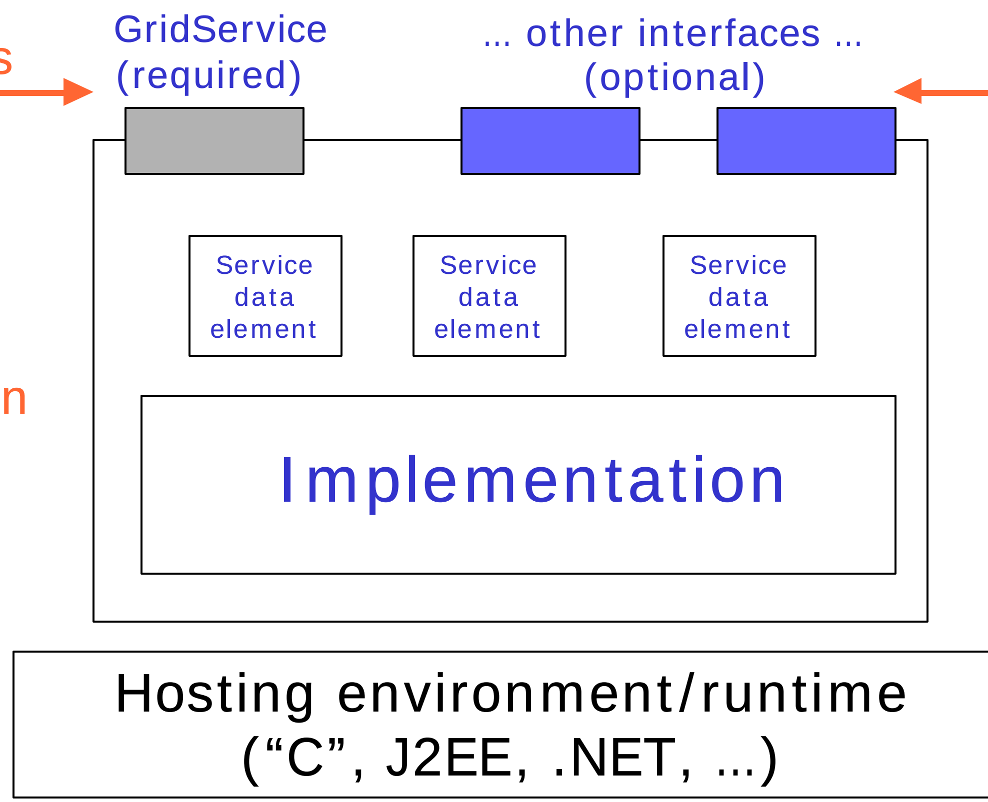

Standard:
- Notification
- Authorization
- Service creation
- Service registry
- Manageability
- Concurrency

+ application-specific interfaces

Service Data

- A Grid service instance maintains a set of service data elements
  - XML fragments encapsulated in standard <name, type, TTL-info> containers
  - Includes basic introspection information, interface-specific data, and application data
- **FindServiceData** operation (GridService interface) queries this information
  - Extensible query language support
- See also notification interfaces
  - Allows notification of service existence and changes in service data

Grid Service Example: Database Service

- A DBaccess Grid service will support at least two portTypes
  - GridService
  - DBaccess

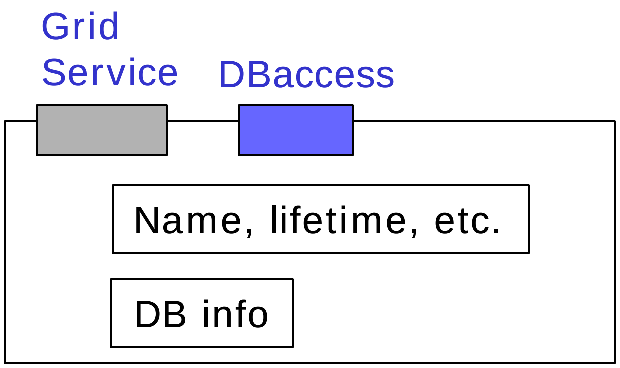

- Each has service data
  - GridService: basic introspection information, lifetime, ...
  - DBaccess: database type, query languages supported, current load, ..., ...
- Maybe other portTypes as well
  - E.g., NotificationSource (SDE = subscribers)

Lifetime Management

- GS instances created by factory or manually; destroyed explicitly or via soft state
  - Negotiation of initial lifetime with a factory (=service supporting Factory interface)
- **GridService** interface supports
  - **Destroy** operation for explicit destruction
  - **SetTerminationTime** operation for keepalive
- Soft state lifetime management avoids
  - Explicit client teardown of complex state
  - Resource “leaks” in hosting environments

Factory

- **Factory** interface's **CreateService** operation creates a new Grid service instance
  - Reliable creation (once-and-only-once)
- **CreateService** operation can be extended to accept service-specific creation parameters
- Returns a **Grid Service Handle (GSH)**
  - A globally unique URL
  - Uniquely identifies the instance for all time
  - Based on name of a home handleMap service

Transient Database Services

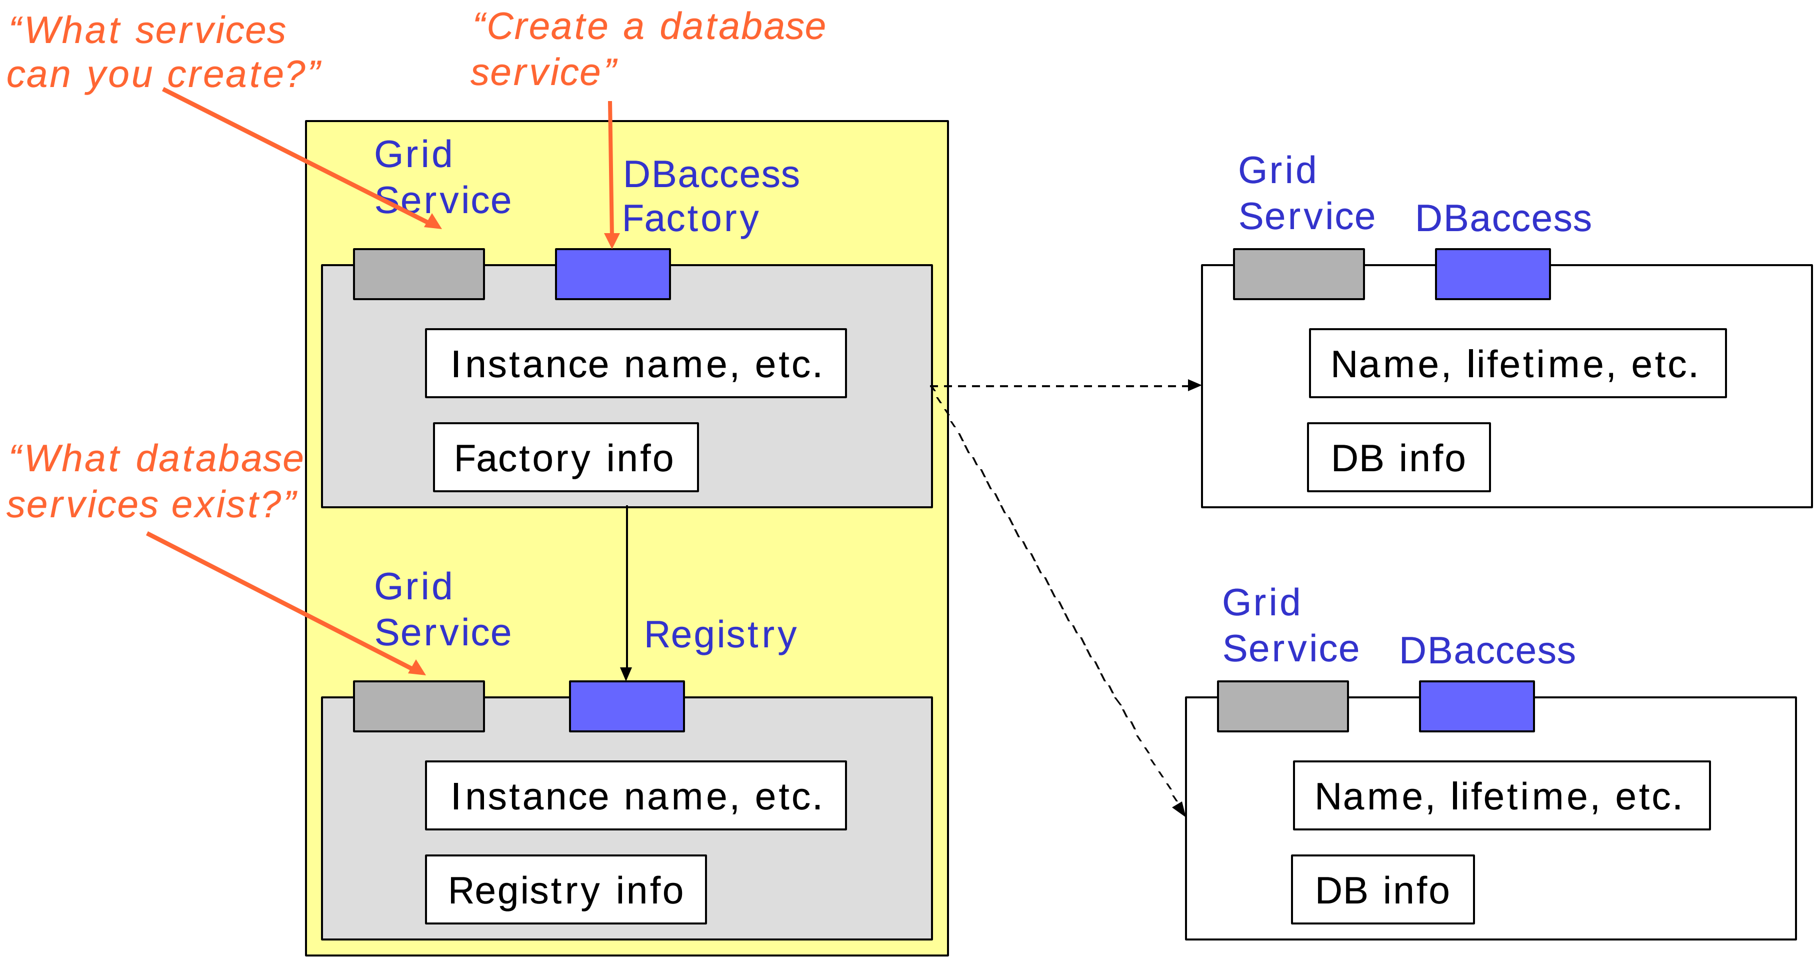

Example: Data Mining for Bioinformatics

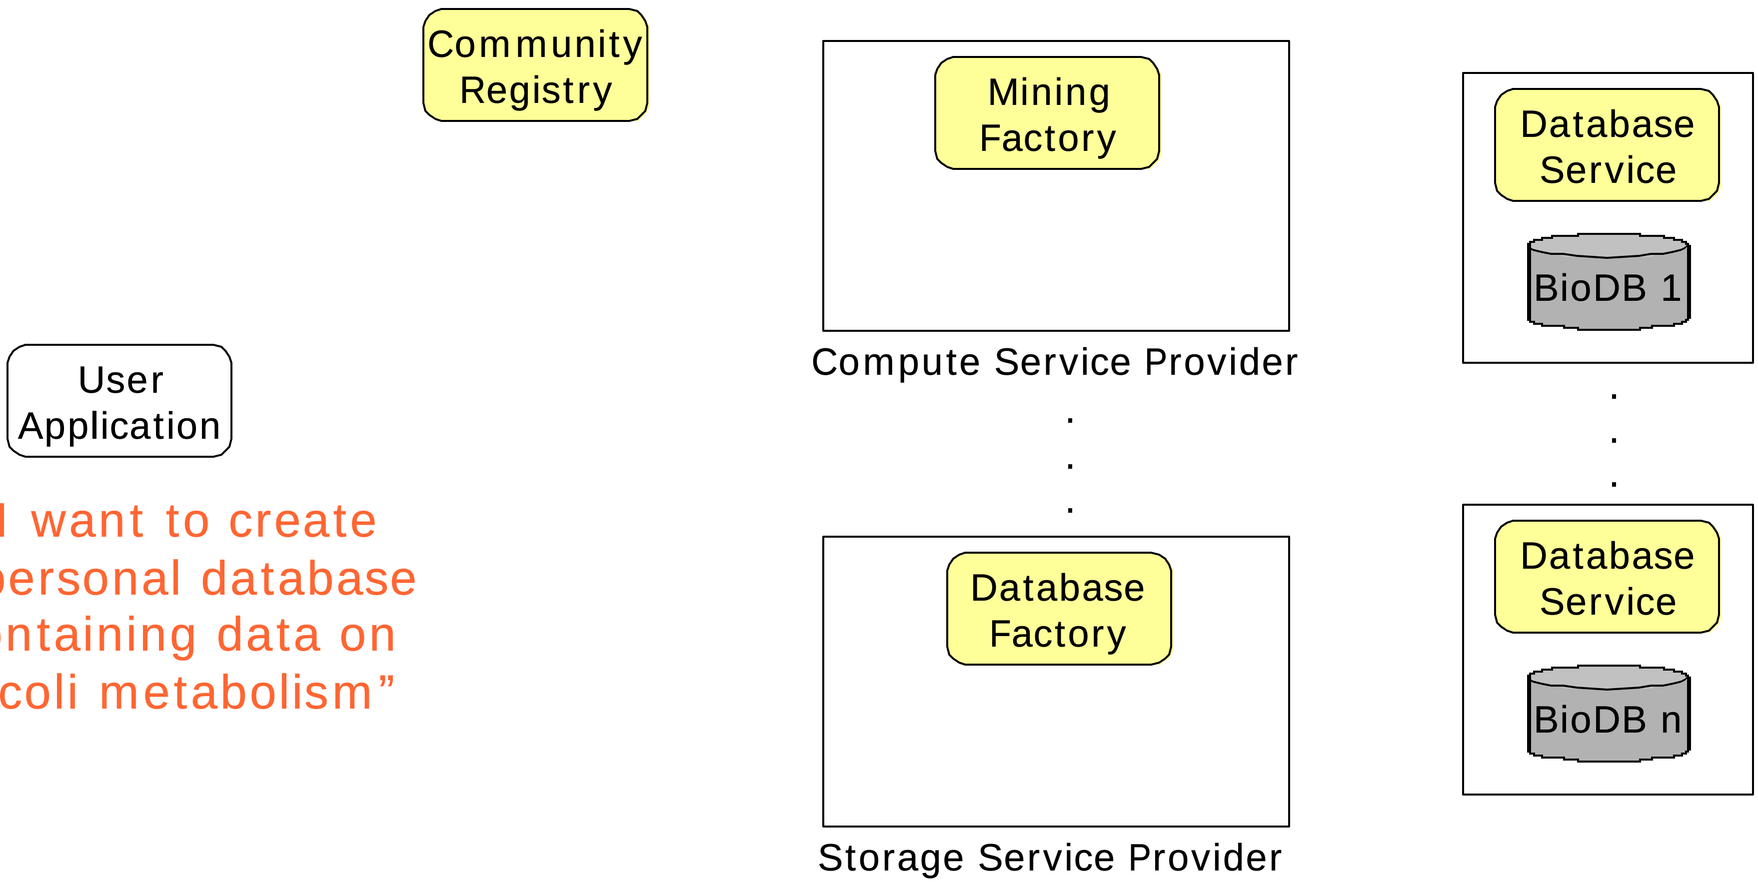

"I want to create a personal database containing data on e.coli metabolism"

Example:
Data Mining for Bioinformatics

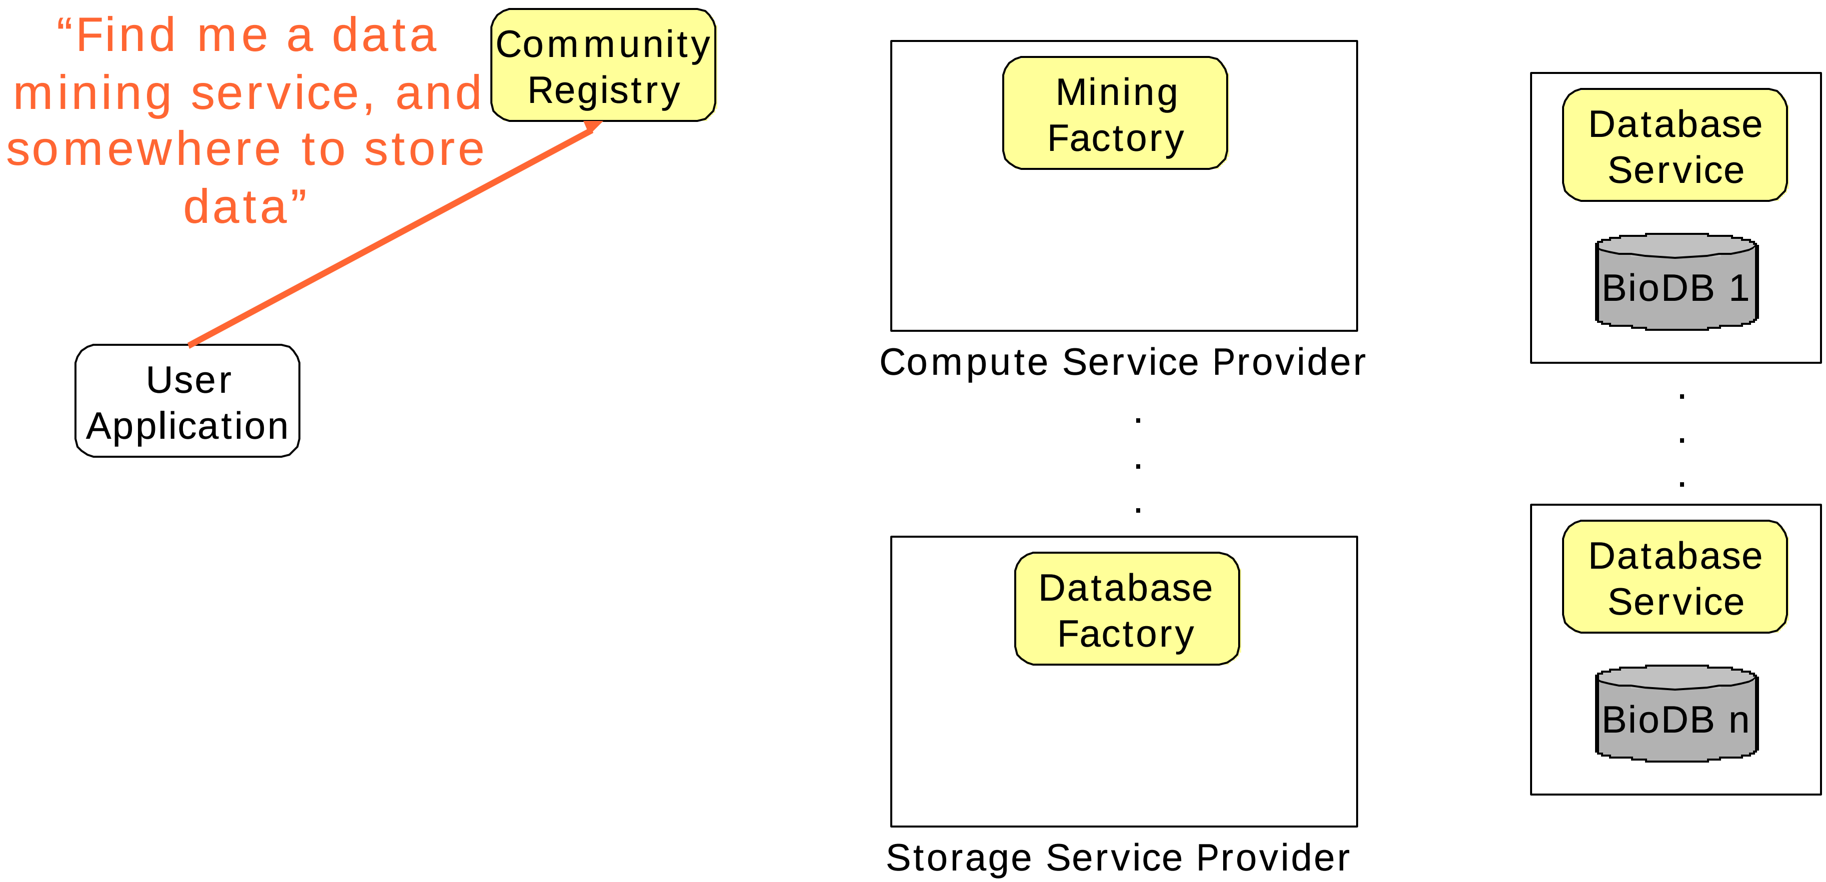

Example:
Data Mining for Bioinformatics

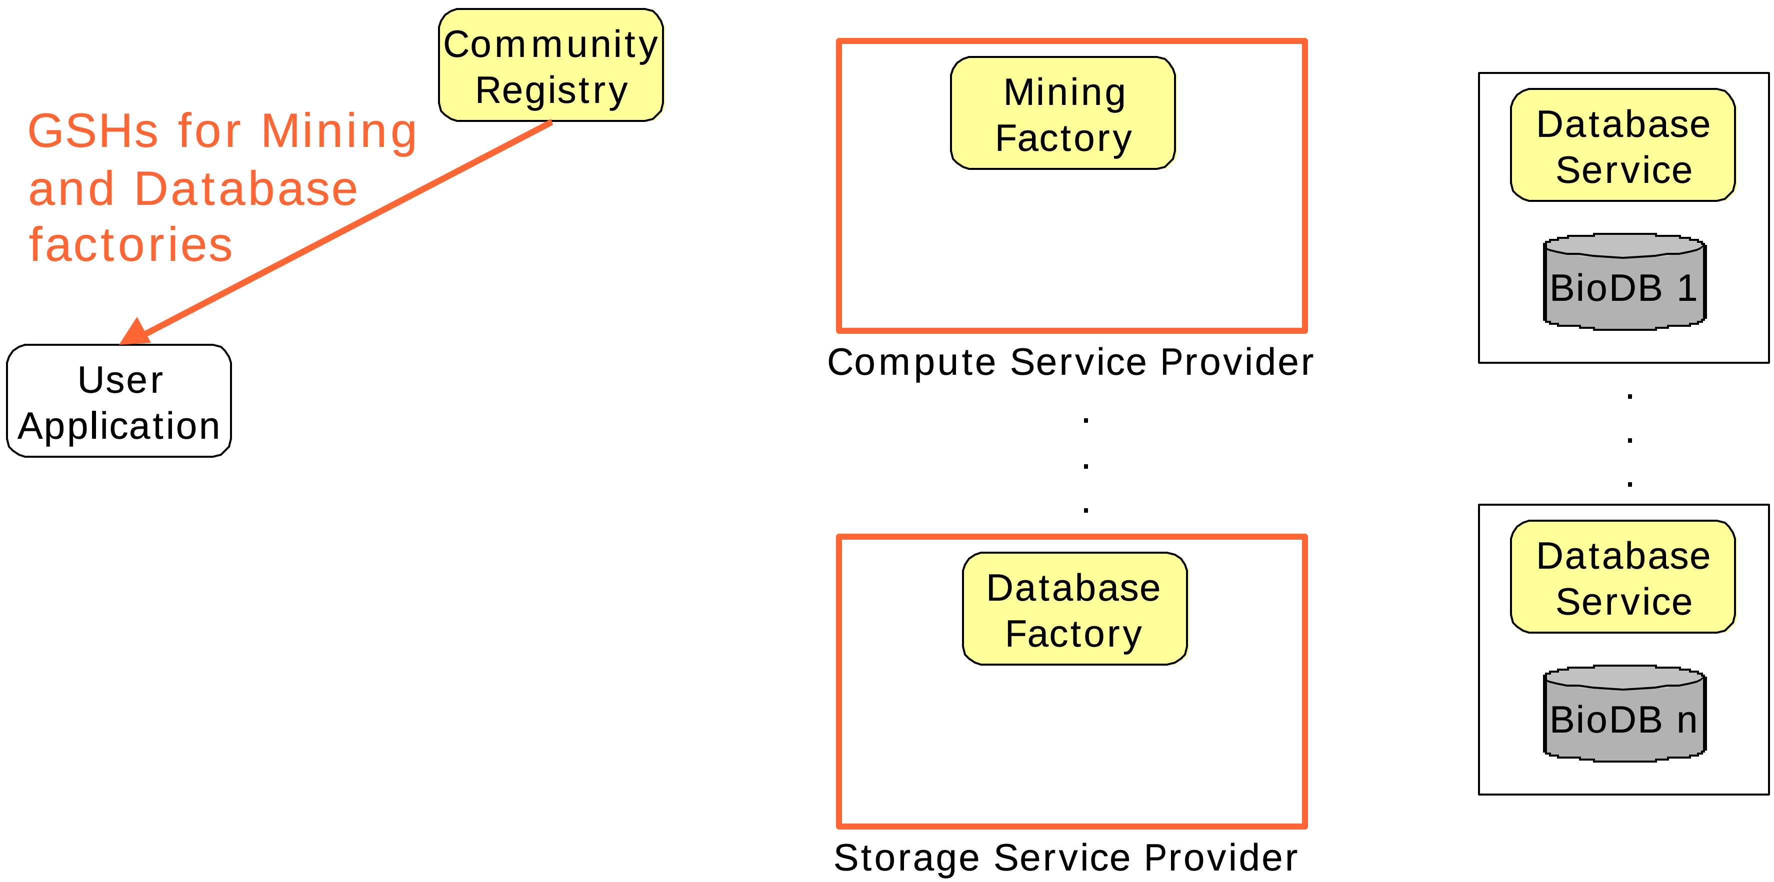

Example: Data Mining for Bioinformatics

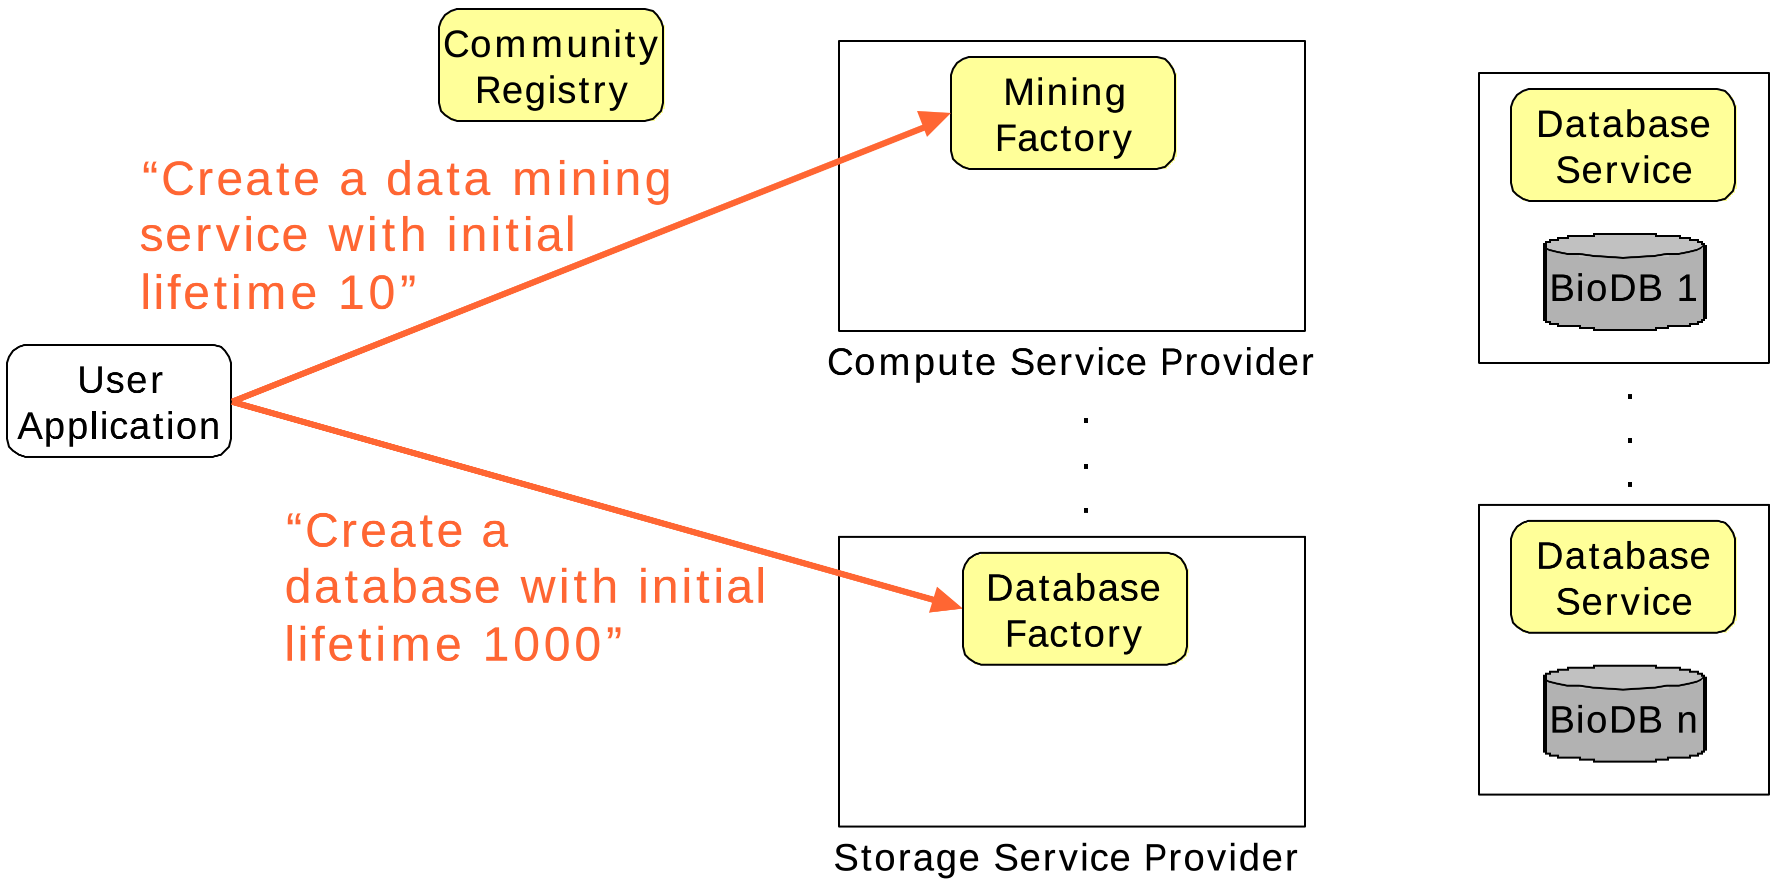

Example: Data Mining for Bioinformatics

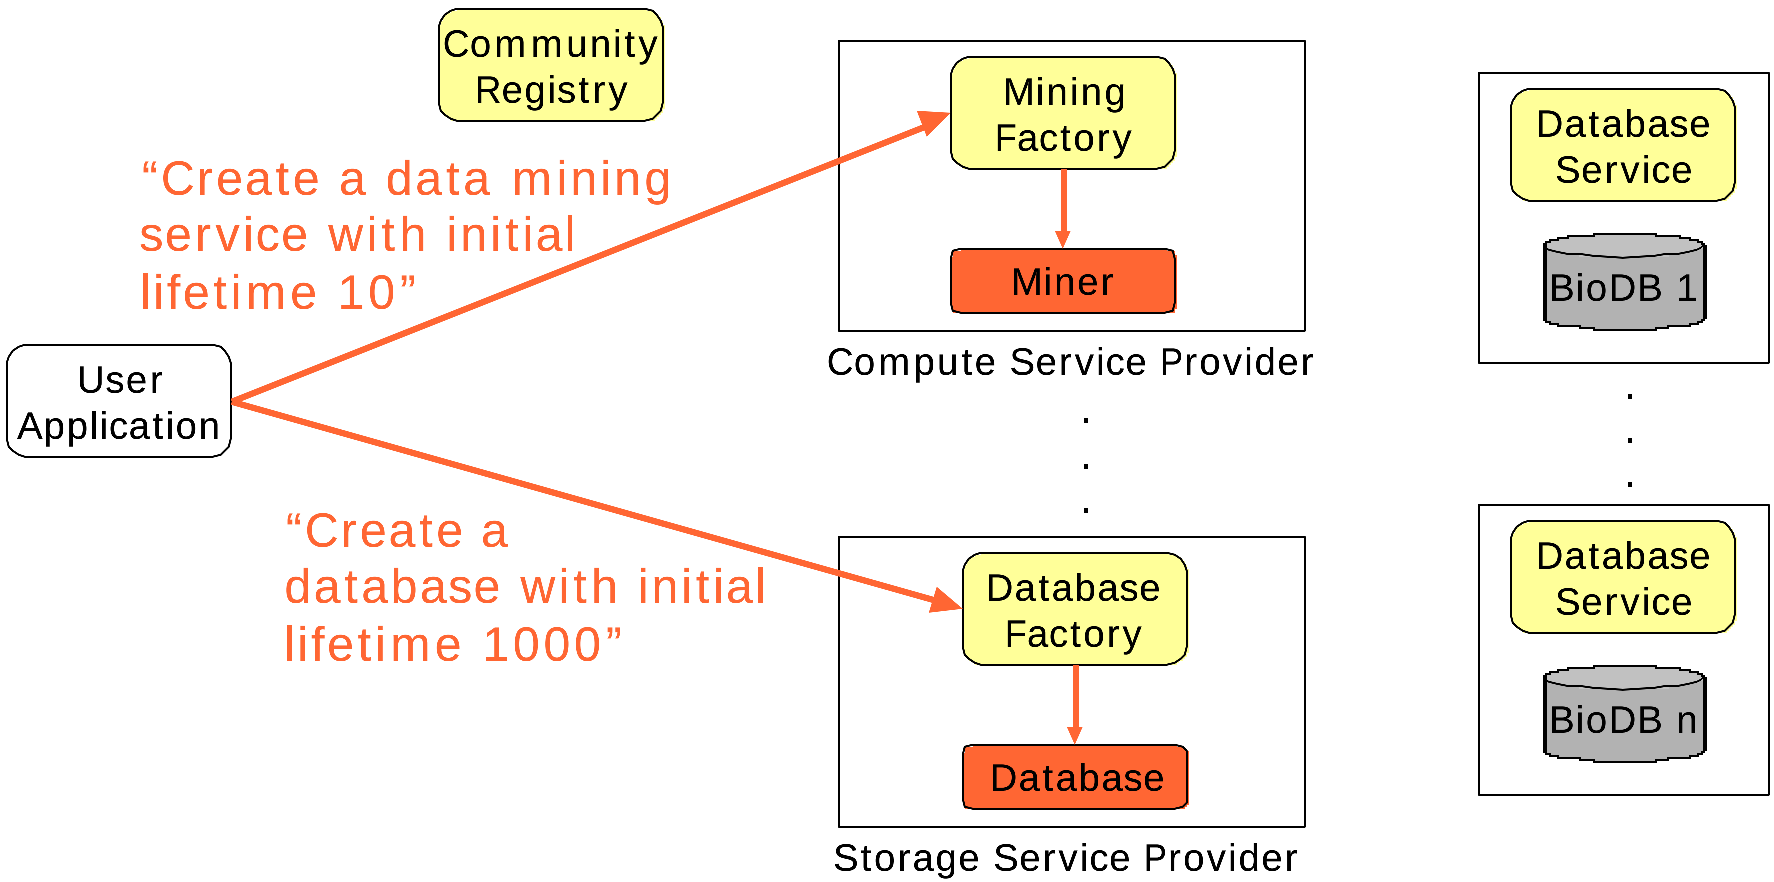

Example: Data Mining for Bioinformatics

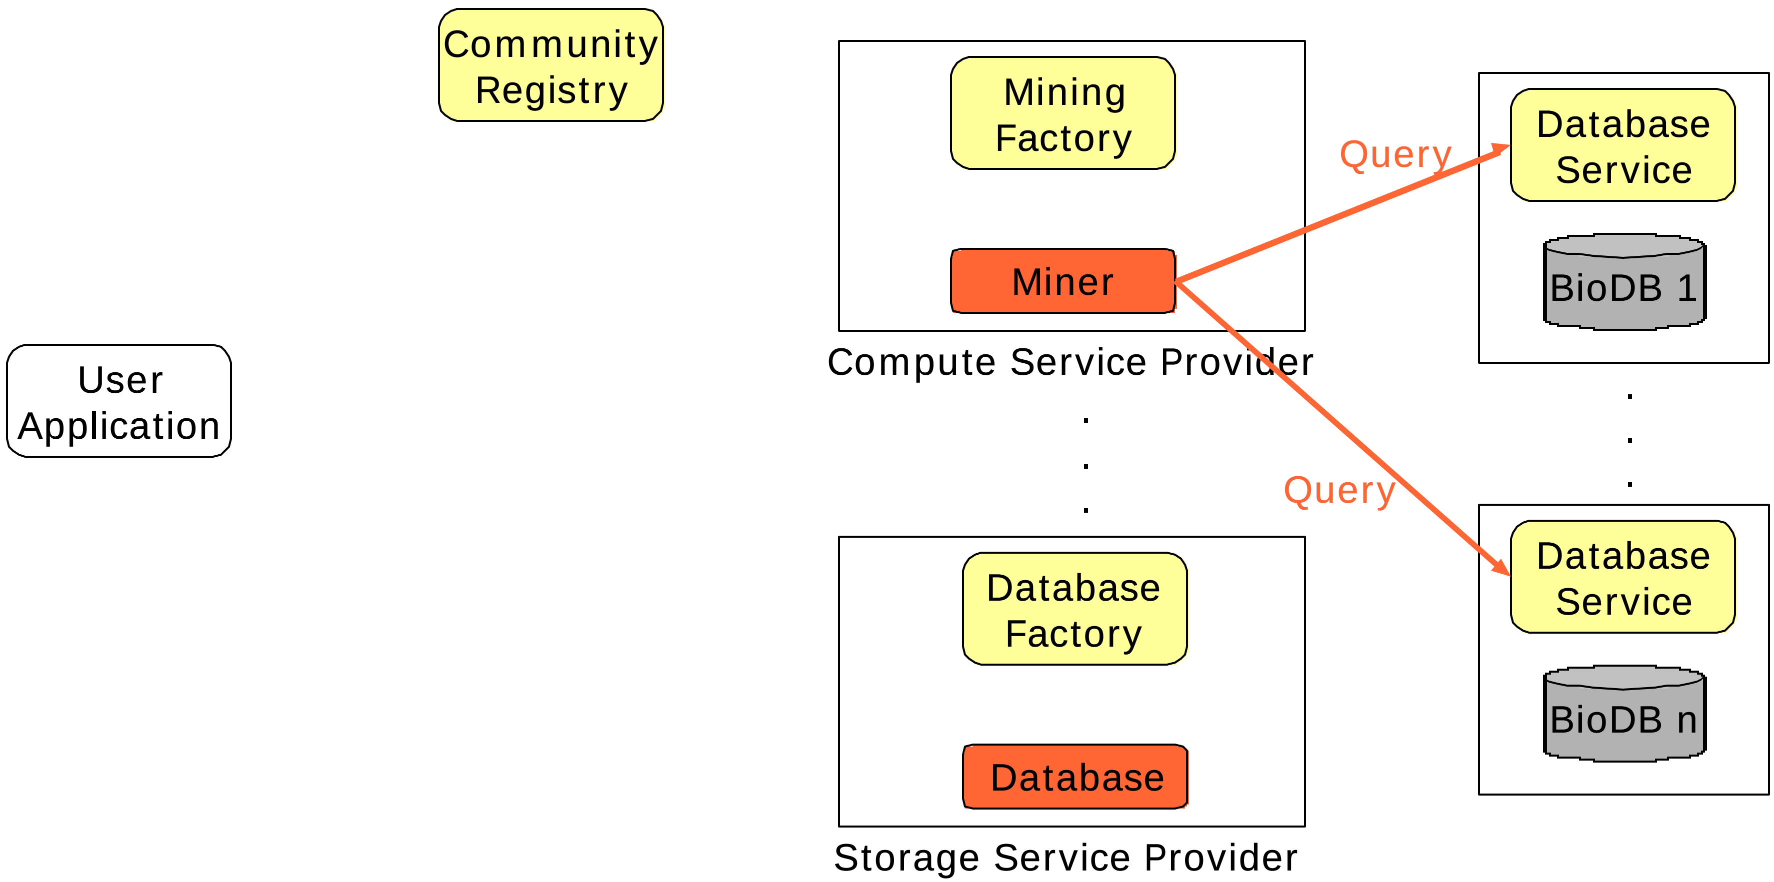

Example:
Data Mining for Bioinformatics

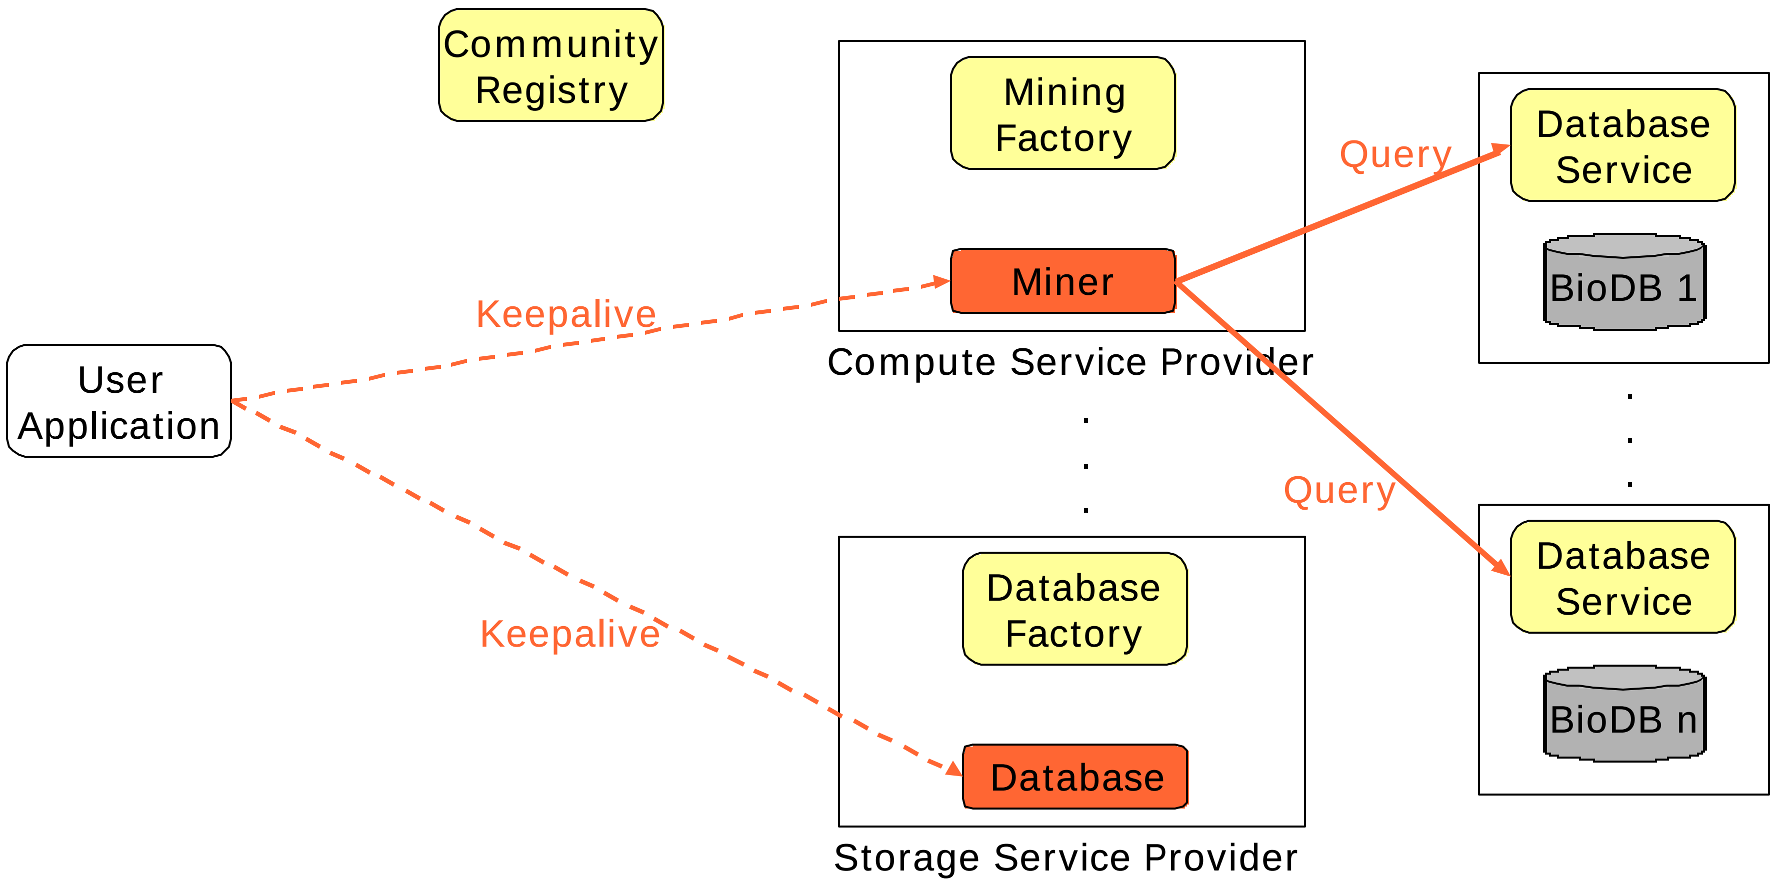

Example: Data Mining for Bioinformatics

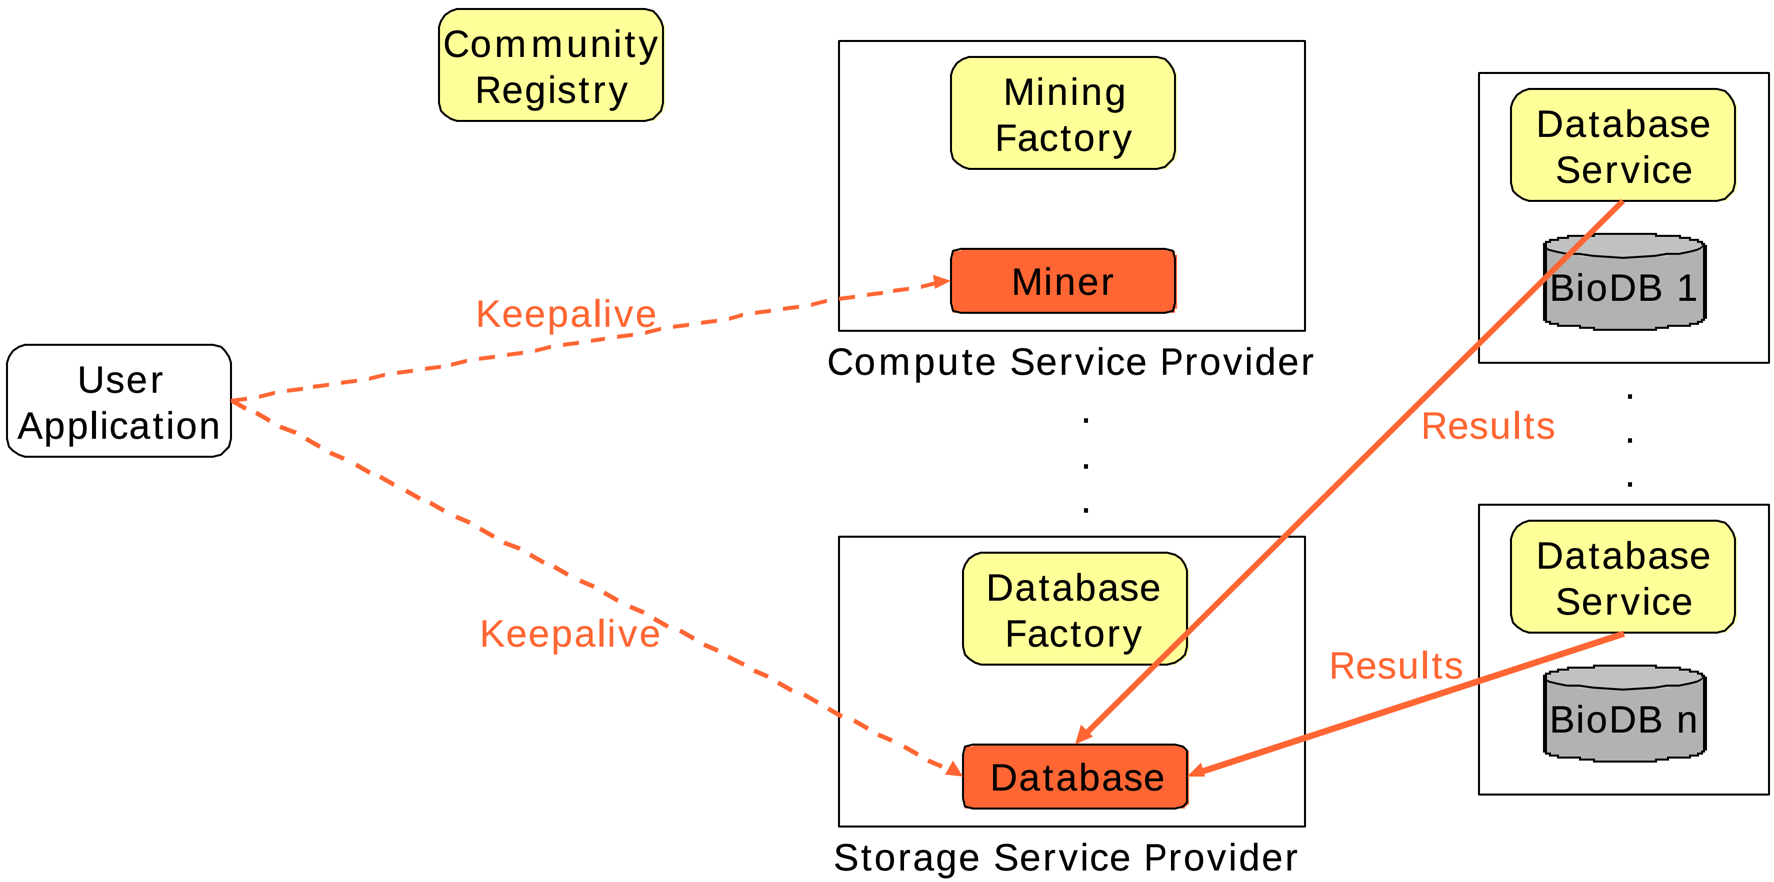

Example:
Data Mining for Bioinformatics

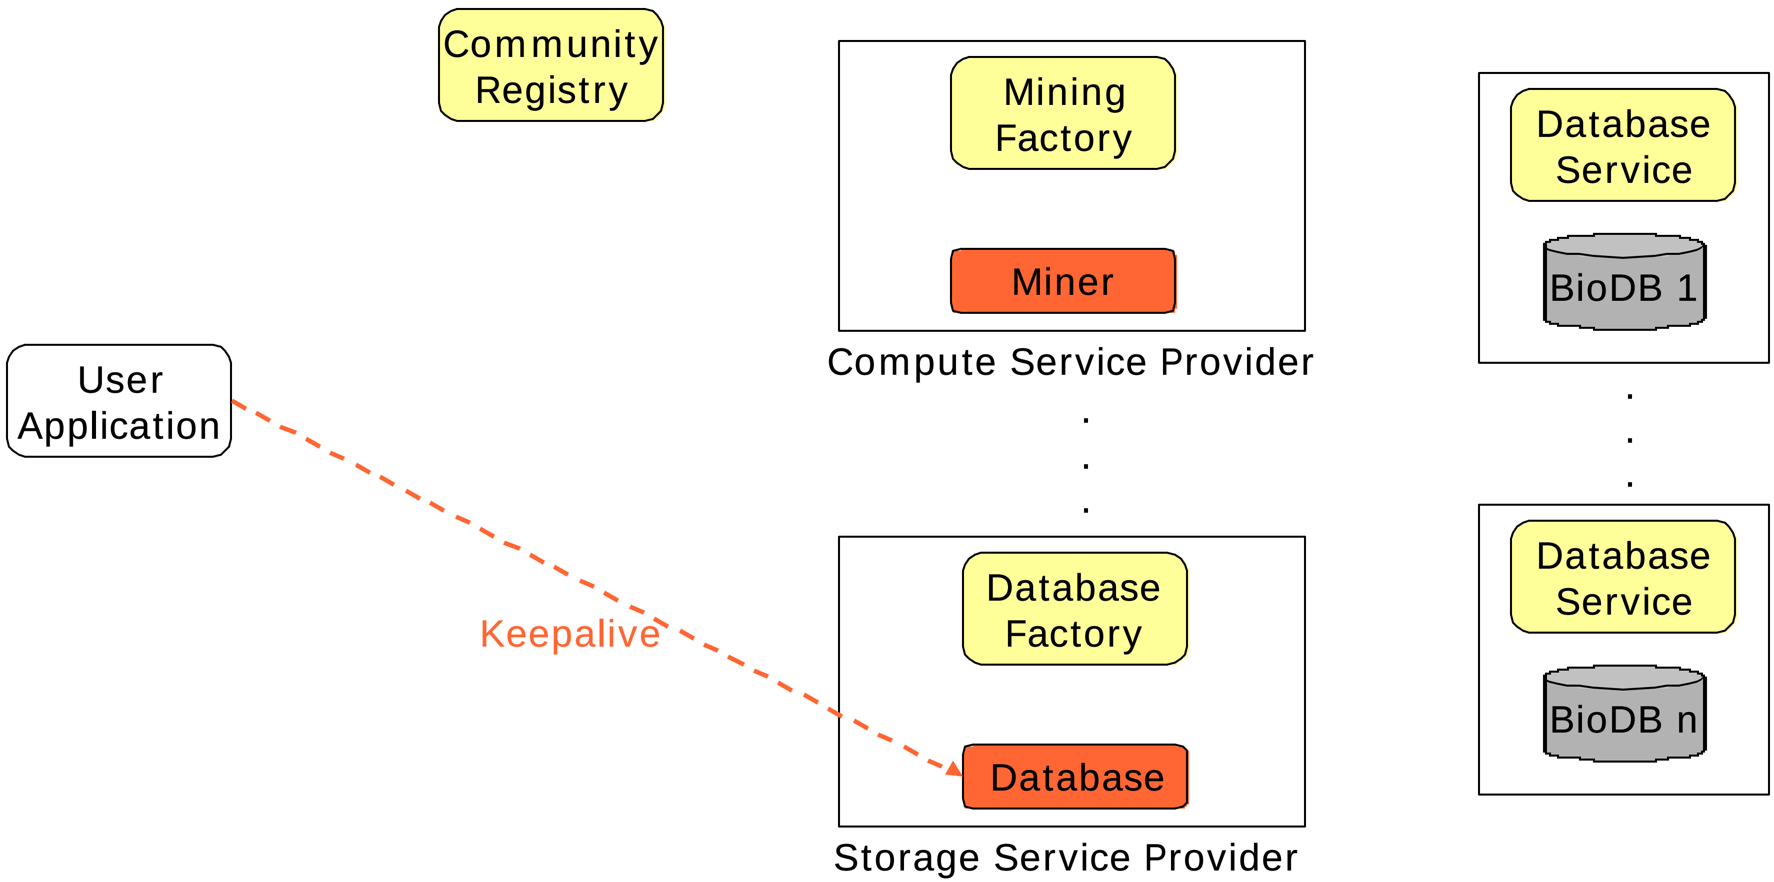

Example:
Data Mining for Bioinformatics

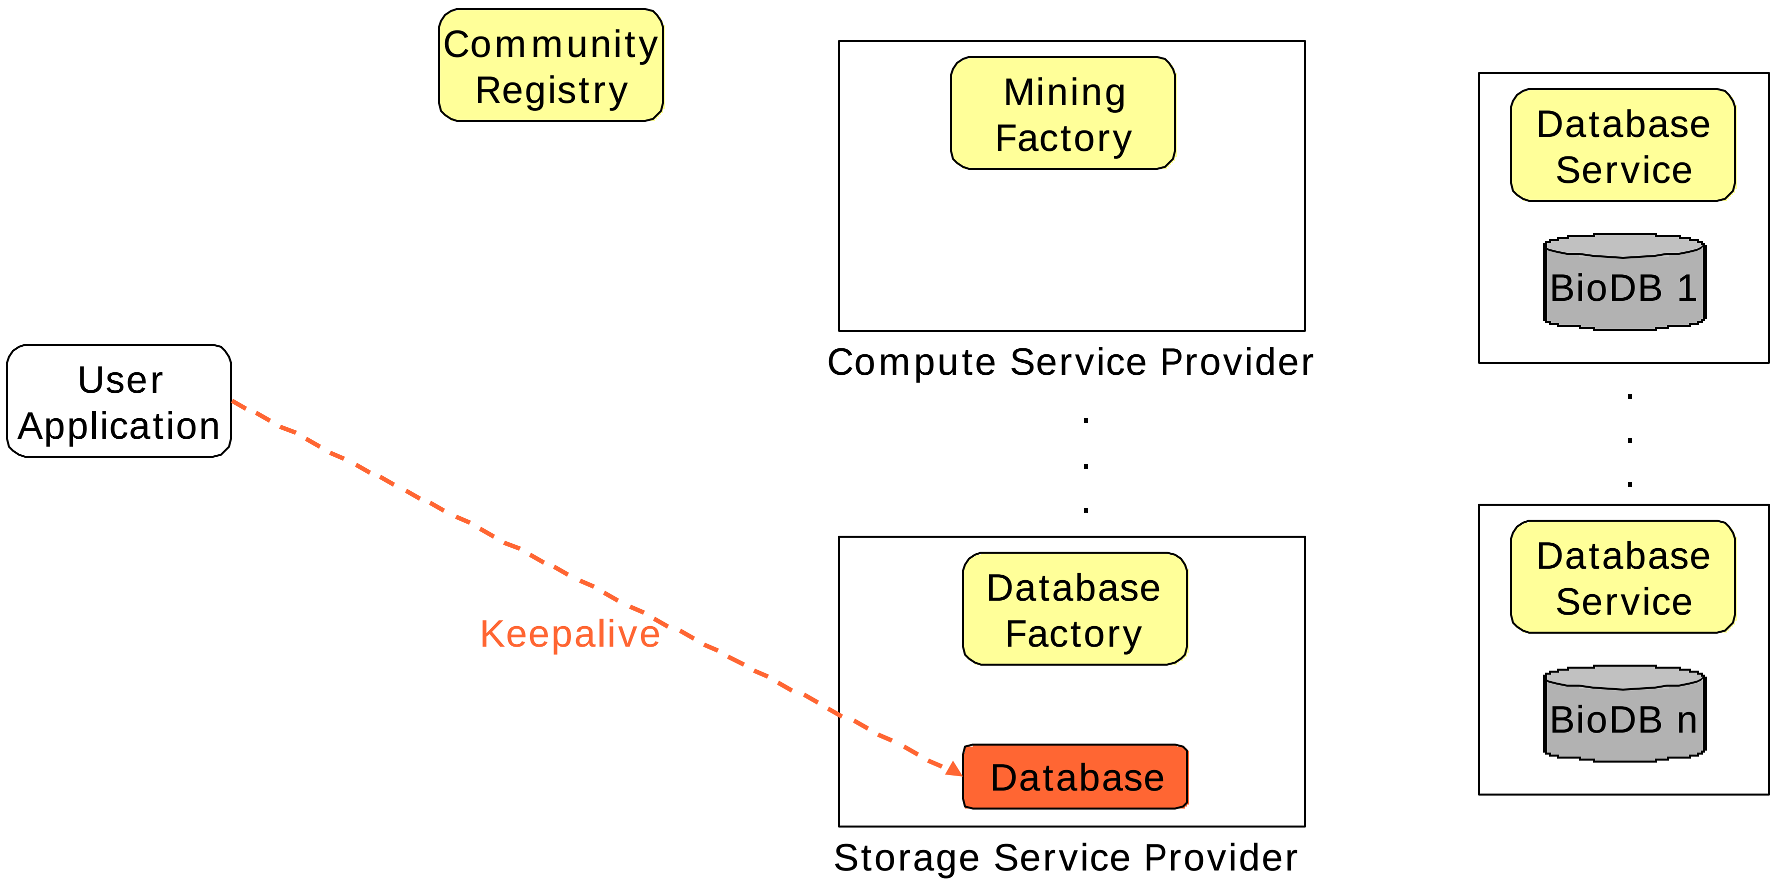

GT3: An Open Source OGSA-Compliant Globus Toolkit

- GT3 Core
  – Implements Grid service interfaces & behaviors
  – Reference impln of evolving standard
  – Multiple hosting envs: Java/J2EE, C, C#/.NET?
- GT3 Base Services
  – Evolution of current Globus Toolkit capabilities
- Many other Grid services

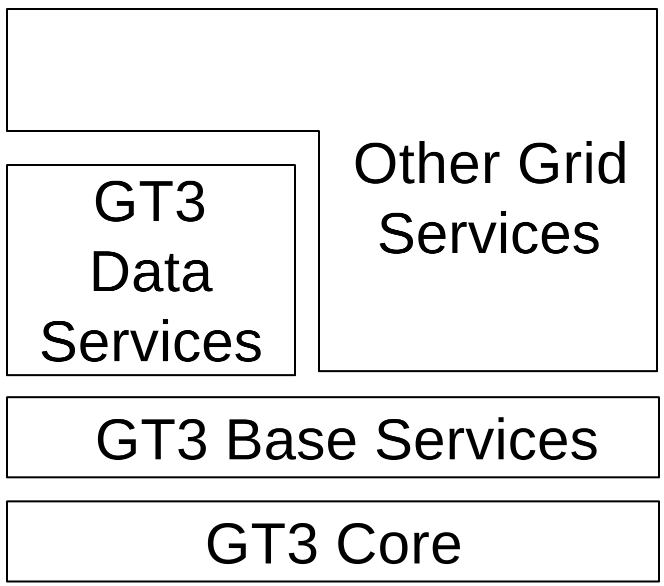

Summary: Grids and Globus Toolkit

- The Grid: Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations
- Considerable impact within eScience, growing interest & adoption within eBusiness
- Globus Toolkit an open source, defacto standard source of protocol and API definitions—and reference implementations
- A strong community organization: the Global Grid Forum

the globus project™
www.globus.org

Summary: Open Grid Services Architecture

- Open Grid Services Architecture represents (we hope!) next step in Grid evolution
- Service orientation enables unified treatment of resources, data, and services
- Standard interfaces and behaviors (the Grid service) for managing distributed state
- Deeply integrated information model for representing and disseminating service data
- Open source Globus Toolkit implementation (and commercial value adds)

For More Information

- Grid Book (somewhat old)
  - www.mkp.com/grids
- Survey + research articles
  - www.mcs.anl.gov/~foster
- The Globus Project™
  - www.globus.org
- GriPhyN project
  - www.griphyn.org
- Global Grid Forum
  - www.gridforum.org
  - www.gridforum.org/ogsi-wg
  - Edinburgh, July 22-24
  - Chicago, Oct 15-17