- **Grid Computing: Resource Sharing and Coordinated Problem Solving in Scalable Distributed Communities**
- Grid computing technologies enable controlled resource sharing in distributed communities and the coordinated use of those shared resources as community members tackle common goals. These technologies include new protocols, services, and APIs for secure resource access, resource management, fault detection, communication, and so forth, that in term enable new application concepts such as virtual data, smart instruments, collaborative design spaces, and metacomputations. In this talk, we review applications that are motivating widespread interest in Grid concepts within the scientific and engineering communities. Then, we describe the Globus Grid architecture that has been adopted by many Grid projects, focusing in particular on our security, resource management, and data management technologies.

# Grid Computing


## Ian Foster

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

## Carl Kesselman

Information Sciences Institute
University of Southern California

# Overview

- Grid computing concept
- Grid technical landscape
- Data grids & data grid projects

# The Grid Concept

# The Grid: The Web on Steroids

# The Grid: The Web on Steroids

Web: Uniform access to HTML documents

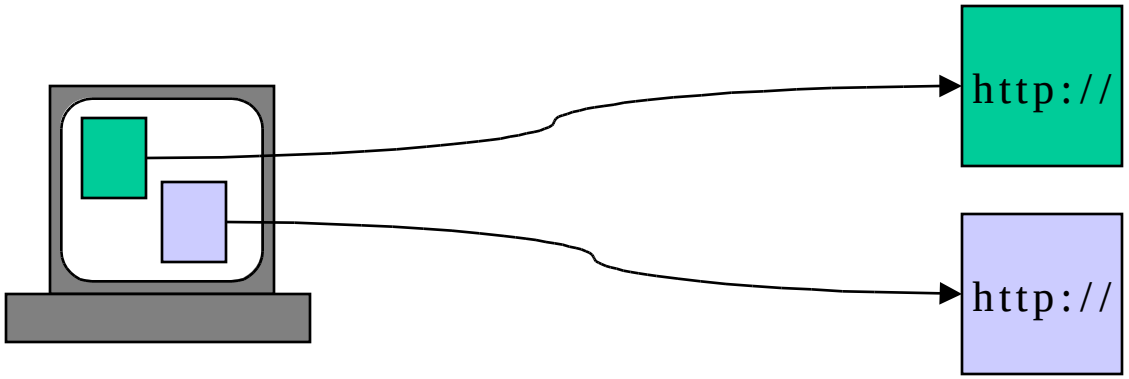

# The Grid: The Web on Steroids

Web: Uniform access to HTML documents

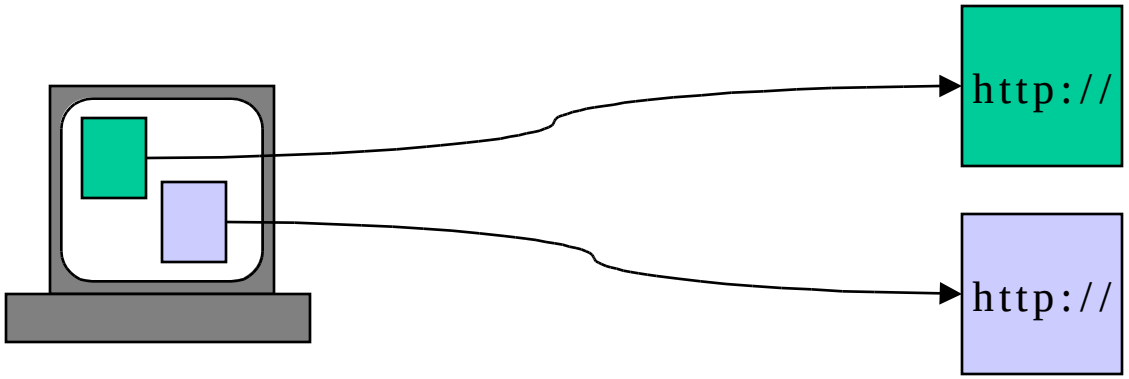


Grid: Flexible, high-perf access to all significant resources

# The Grid: The Web on Steroids

Web: Uniform access to HTML documents

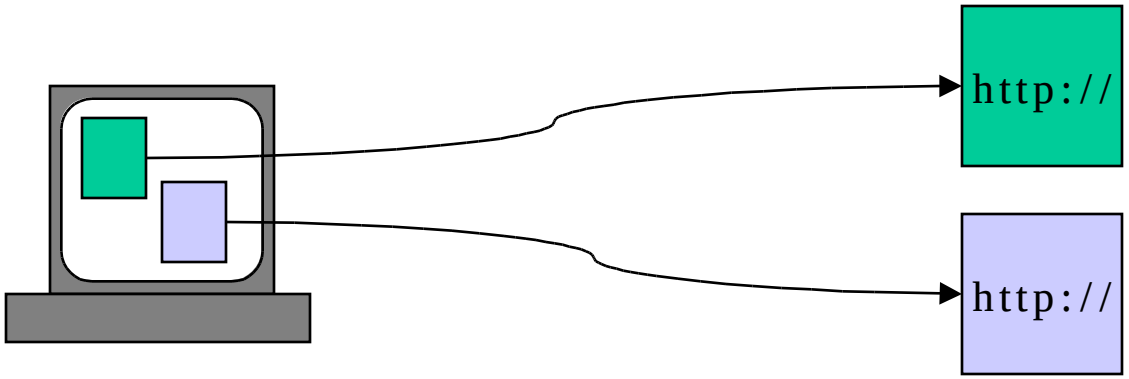


Grid: Flexible, high-perf access to all significant resources

# The Grid: The Web on Steroids

Web: Uniform access to HTML documents

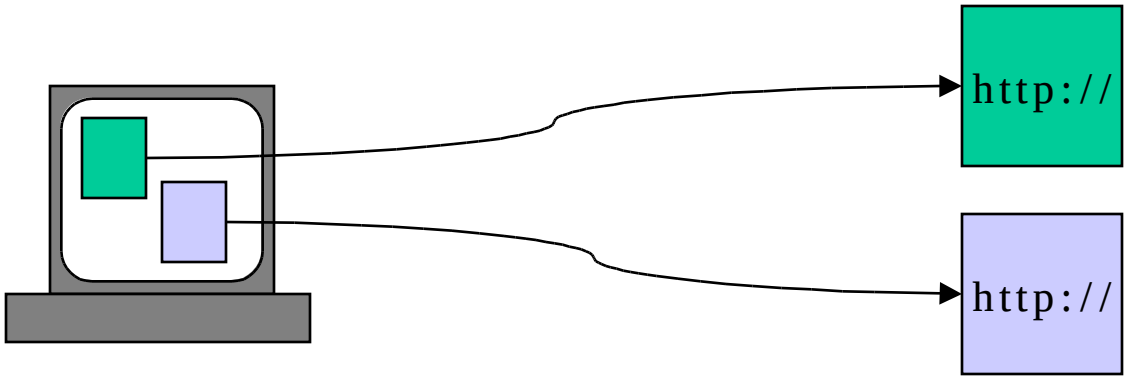


Grid: Flexible, high-perf access to all significant resources

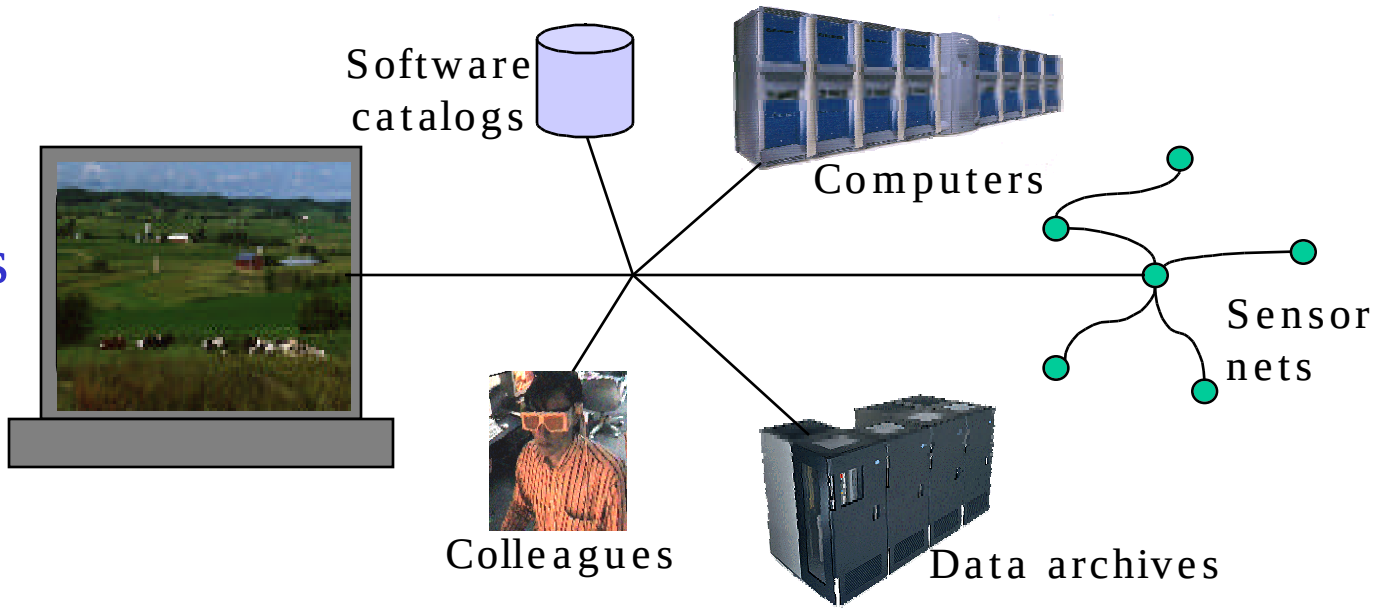

# The Grid: The Web on Steroids

Web: Uniform access to HTML documents

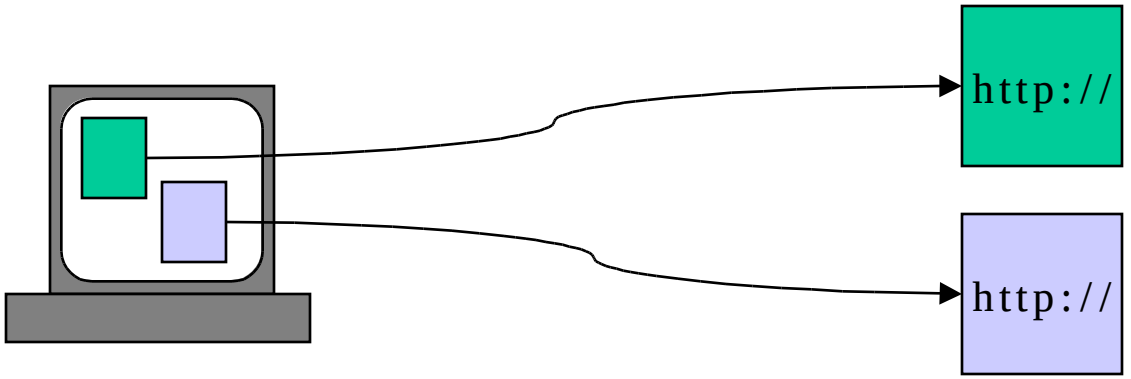


Grid: Flexible, high-perf access to all significant resources

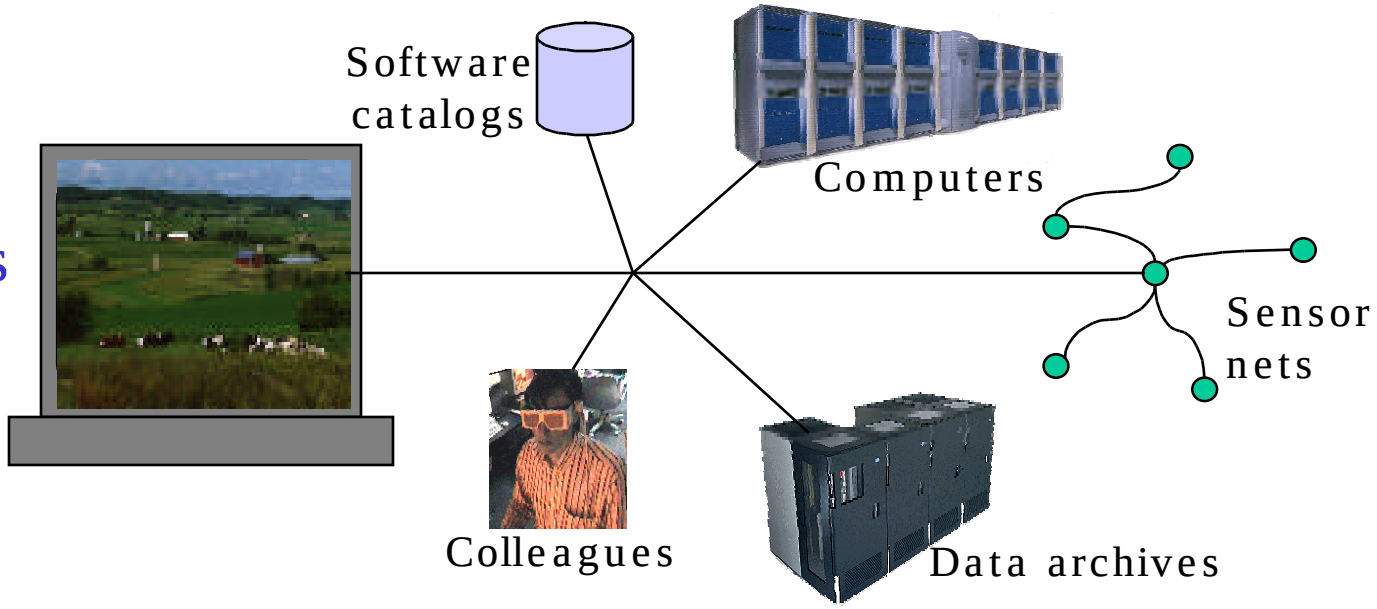


*On-demand creation of powerful virtual computing systems*

foster@mcs.anl.gov

# Grid Computing: Take 2

- Enable communities (“virtual organizations”) to share geographically distributed resources as they pursue common goals—in the absence of central control, omniscience, trust relationships
- Via investigations of
  - New applications that become possible when resources can be shared in a coordinated way
  - Protocols, algorithms, persistent infrastructure to facilitate sharing

# Grid Communities and Applications: NSF National Technology Grid

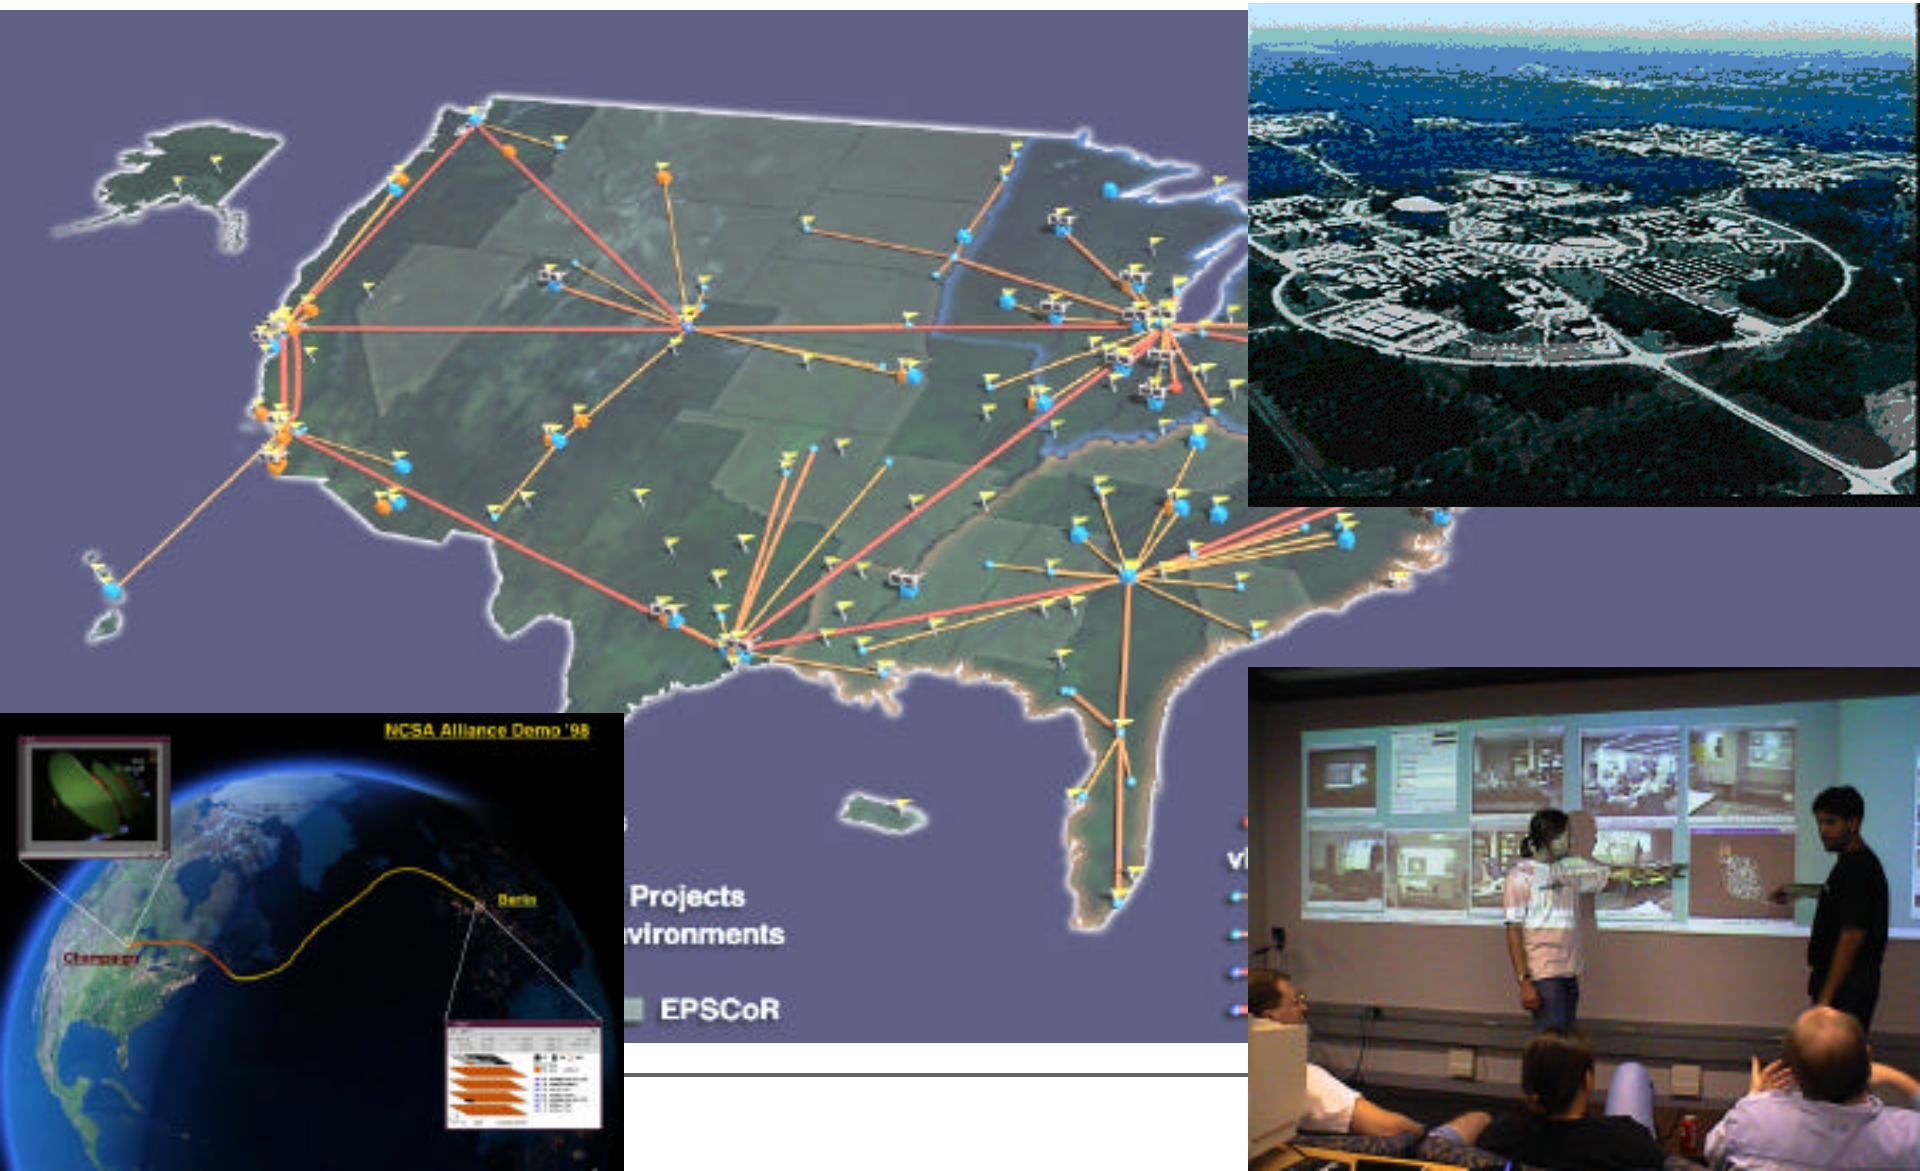

# Grid Communities and Applications: Network for Earthquake Eng. Simulation

- NEESgrid: national infrastructure to couple earthquake engineers with experimental facilities, databases, computers, & each other
- On-demand access to experiments, data streams, computing, archives, collaboration

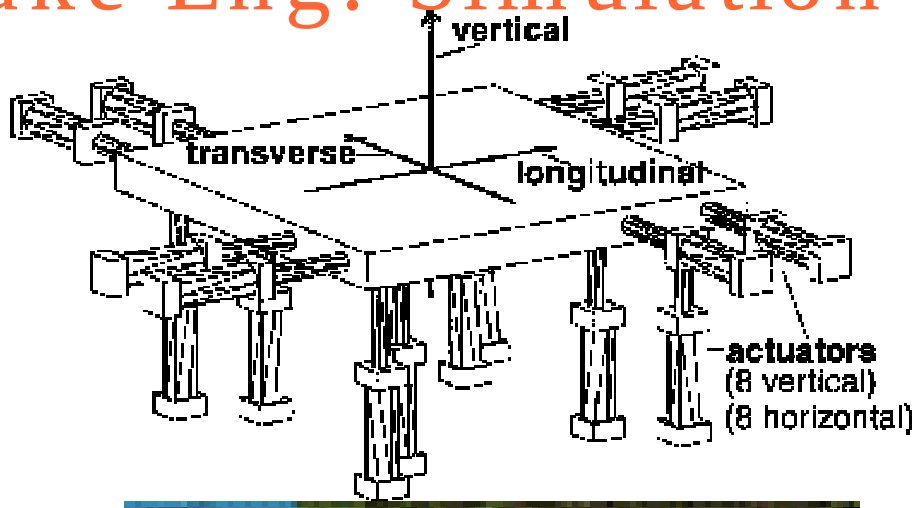


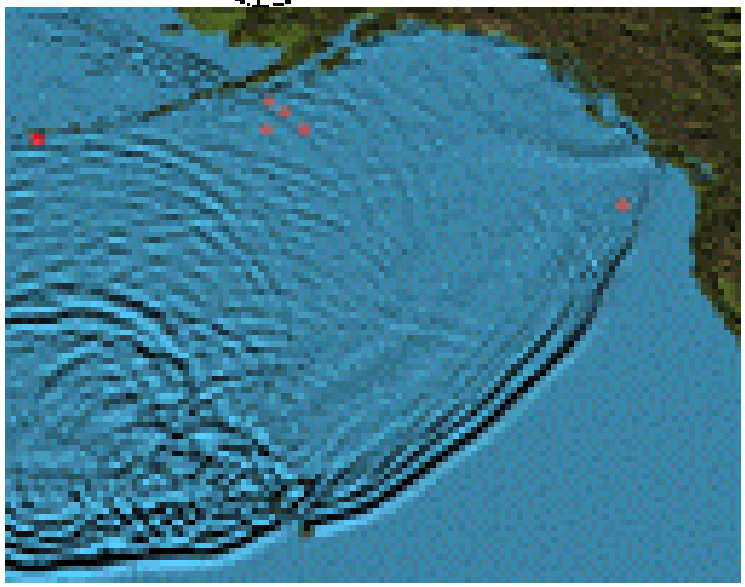

NEESgrid: Argonne, Michigan, NCSA, UIUC, USC

# Grid Communities and Applications: Mathematicians Solve NUG30

- Community = an informal collaboration of mathematicians and computer scientists
- Condor-G delivers 3.46E8 CPU seconds in 7 days (peak 1009 processors) in U.S. and Italy (8 sites)
- Solves NUG30 quadratic assignment problem

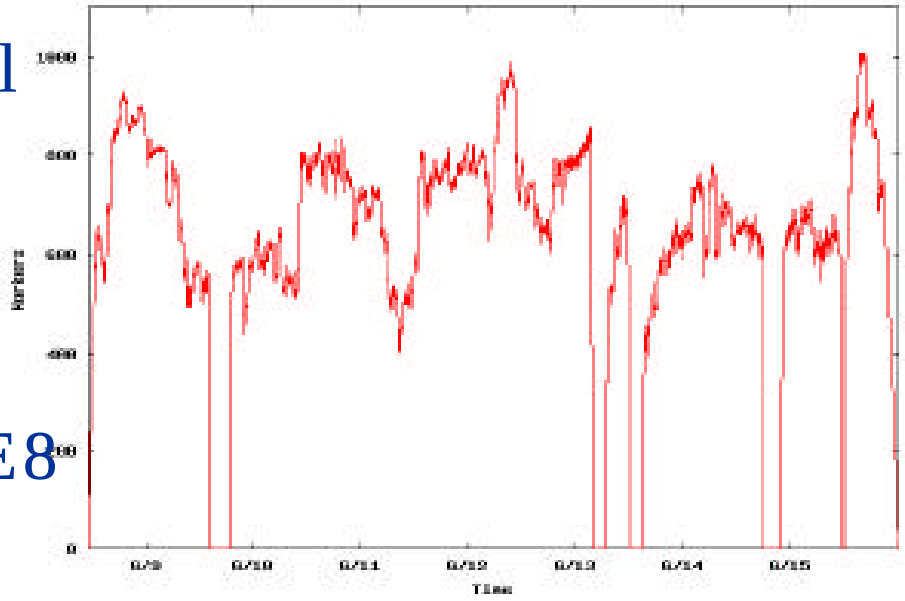


14,5,28,24,1,3,16,15,
10,9,21,2,4,29,25,22,
13,26,17,30,6,20,19,
8,18,7,27,12,11,23

MetaNEOS: Argonne, Iowa, Northwestern, Wisconsin

# Grid Communities and Applications: Home Computers Evaluate AIDS Drugs

- Community =
  - 1000s of home computer users
  - Philanthropic computing vendor (Entropia)
  - Research group (Scripps)
- Common goal= advance AIDS research

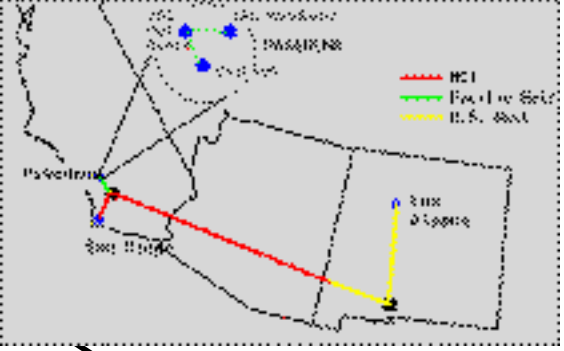

# A Little History

- Early 90s
  - Gigabit testbeds, metacomputing
- Mid to late 90s
  - Early experiments (e.g., I-WAY), academic software projects (e.g., Globus), application experiments
- 2000
  - Major application communities emerging
  - Major infrastructure deployments
  - Rich technology base
  - Grid Forum: >300 people, >90 orgs, 11 countries

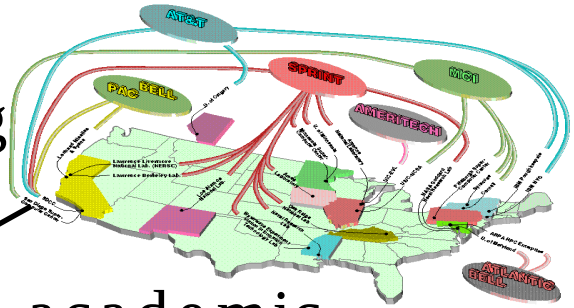

# Technical Landscape

# Grid Technical Challenges: Sharing, Coordinated Use

- New problem solving methods and tools
  - Data, computing, collaboration, sensors, ...
- Numerous cross-cutting technical problems
  - Authentication, authorization, policy, audit; resource discovery, access, allocation, control; failure detection & recovery; brokering; ...; ...
- No central control, omniscience, trust; hence
  - Need to preserve local site independence
  - Policy discovery and negotiation important
  - Myriad interesting failure modes

# How Are Cross-Cutting Technical Issues Addressed?

- Development of Grid protocols & services
  - Protocol-mediated access to remote resources
  - New services: e.g., resource brokering
  - "On the Grid" = speak Intergrid protocols
  - Mostly (extensions to) existing protocols
- Development of Grid APIs & SDKs
  - Facilitate application development by supplying higher-level abstractions
- The (hugely successful) model is the Internet
- The Grid is not a distributed OS!

# Layered Grid Architecture (By Analogy to Internet Architecture)

# Layered Grid Architecture (By Analogy to Internet Architecture)

"Controlling things locally": Access to, & control of, resources

Fabric

# Layered Grid Architecture (By Analogy to Internet Architecture)

| | |
|---|---|
| "Talking to things": communication (Internet protocols) & security | Connectivity |
| "Controlling things locally": Access to, & control of, resources | Fabric |

# Layered Grid Architecture (By Analogy to Internet Architecture)

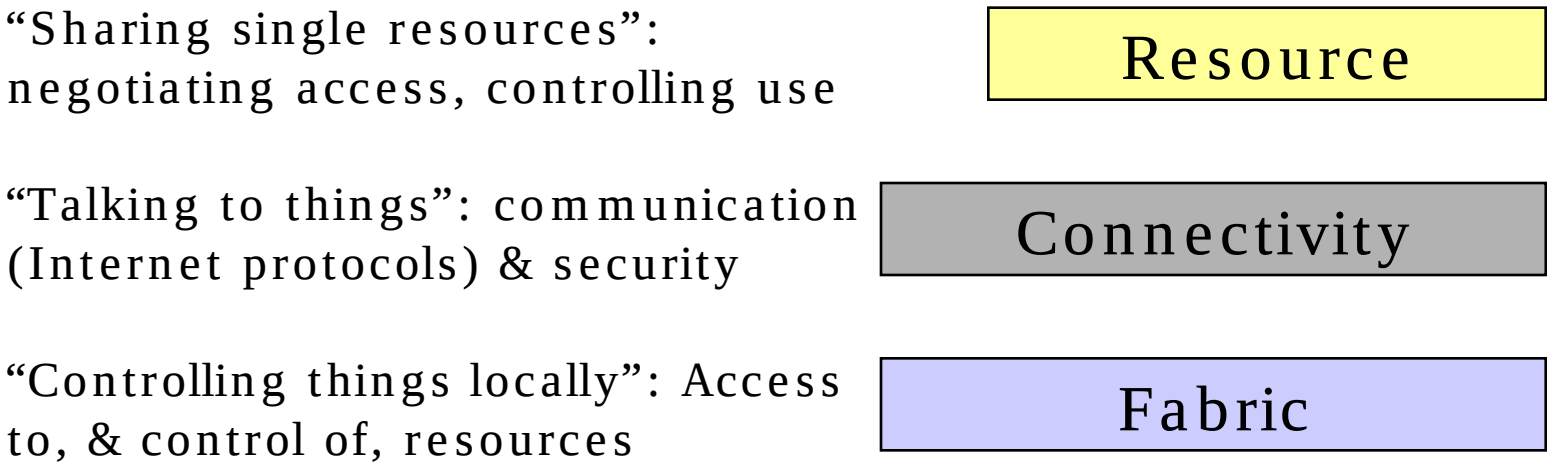

# Layered Grid Architecture (By Analogy to Internet Architecture)

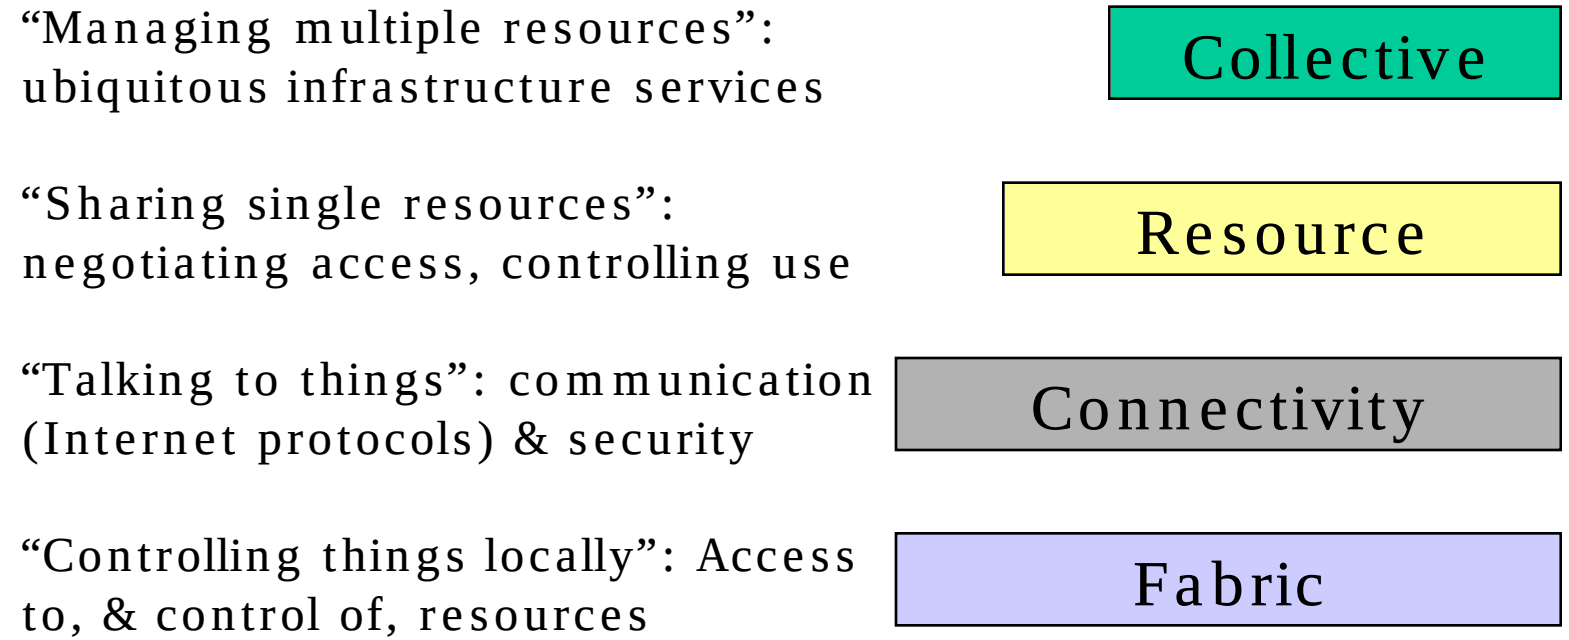

# Layered Grid Architecture (By Analogy to Internet Architecture)

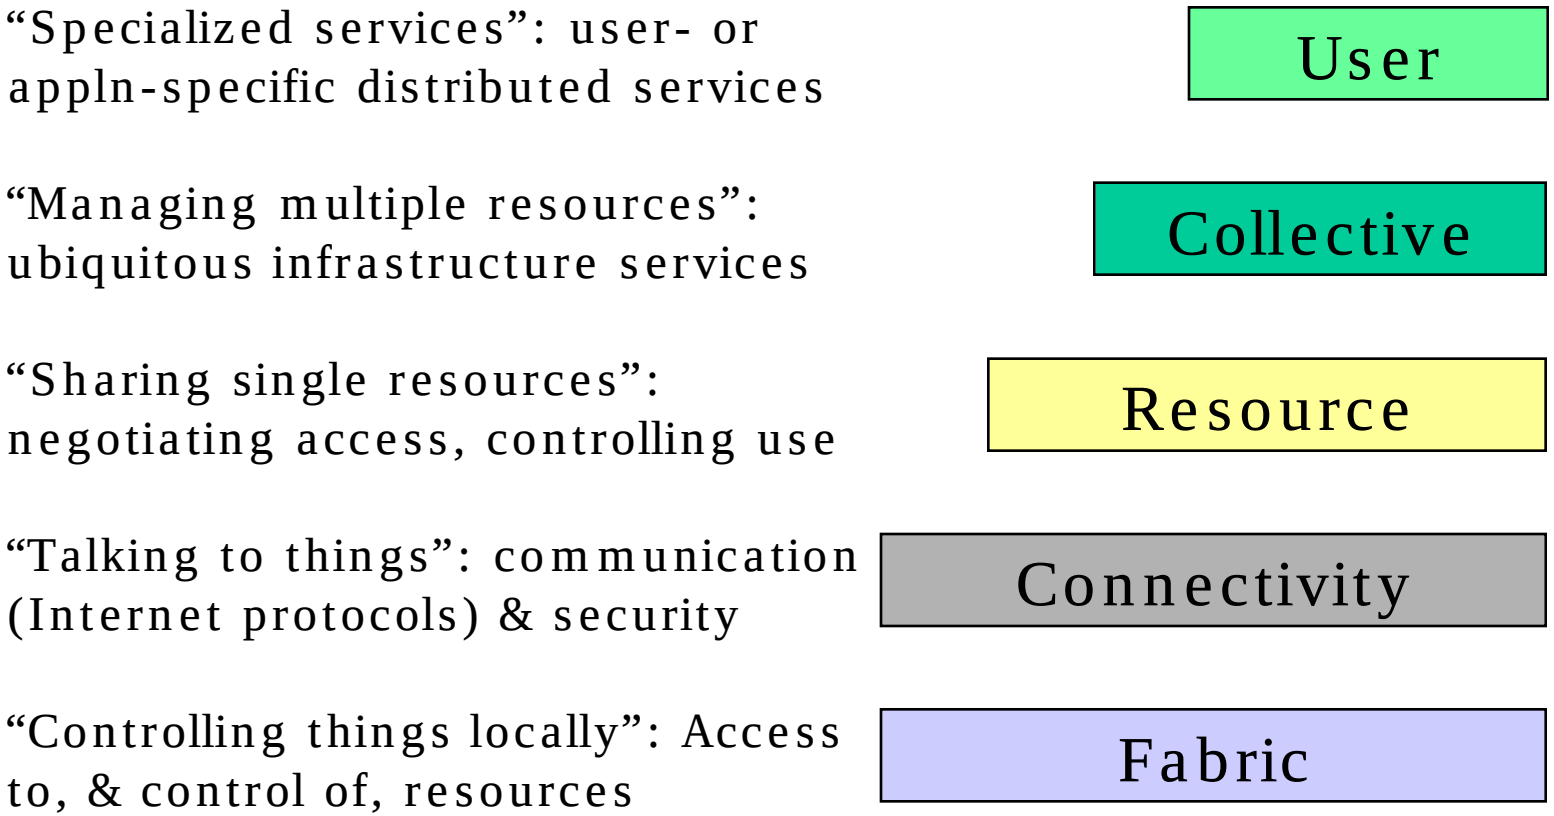

# Layered Grid Architecture (By Analogy to Internet Architecture)

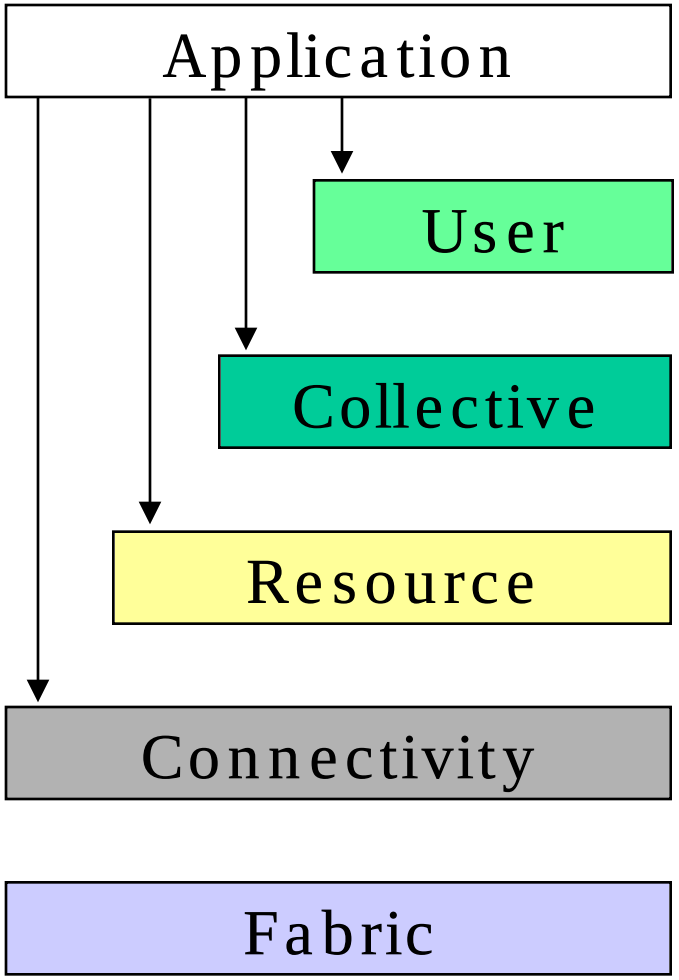


"Specialized services": user- or appln-specific distributed services

"Managing multiple resources": ubiquitous infrastructure services

"Sharing single resources": negotiating access, controlling use

"Talking to things": communication (Internet protocols) & security

"Controlling things locally": Access to, & control of, resources

# Layered Grid Architecture (By Analogy to Internet Architecture)

"Specialized services": user- or appln-specific distributed services

"Managing multiple resources": ubiquitous infrastructure services

"Sharing single resources": negotiating access, controlling use

"Talking to things": communication (Internet protocols) & security

"Controlling things locally": Access to, & control of, resources

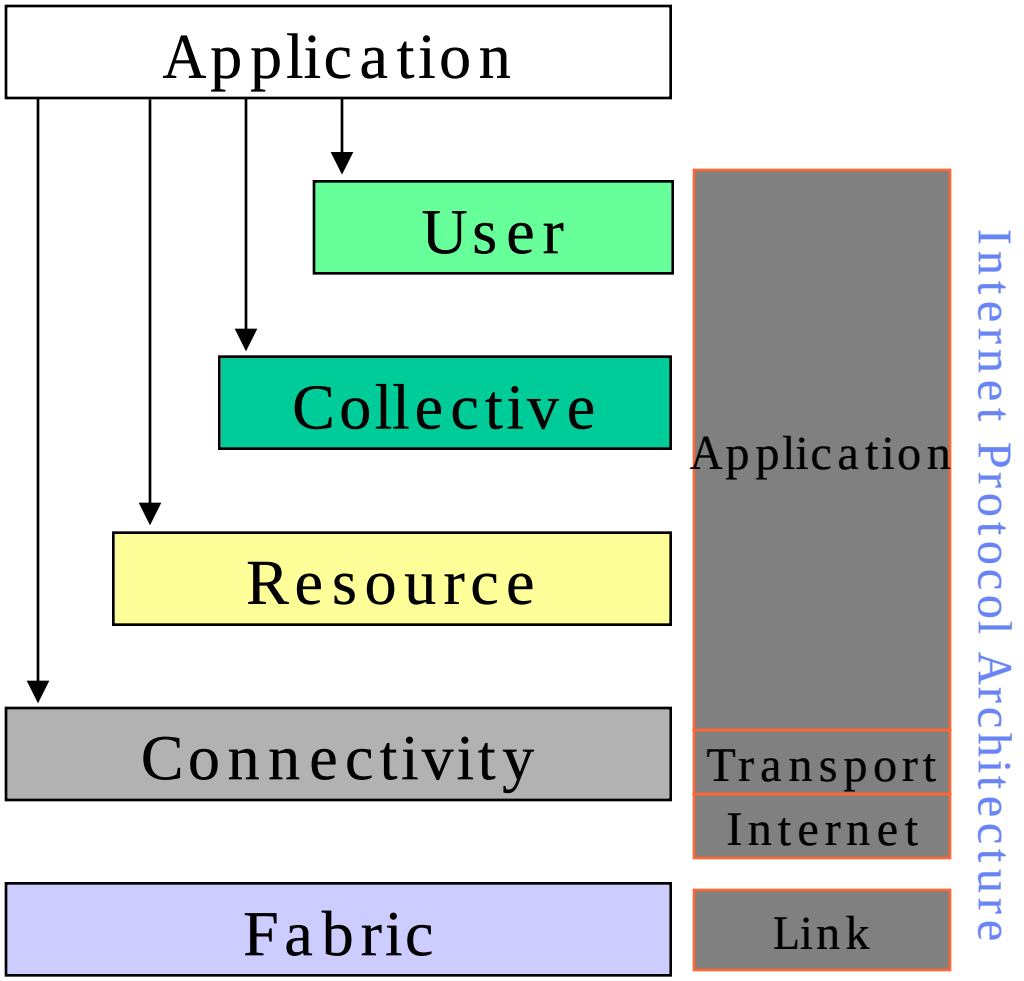

# Protocols, Services, and Interfaces Occur at Each Level

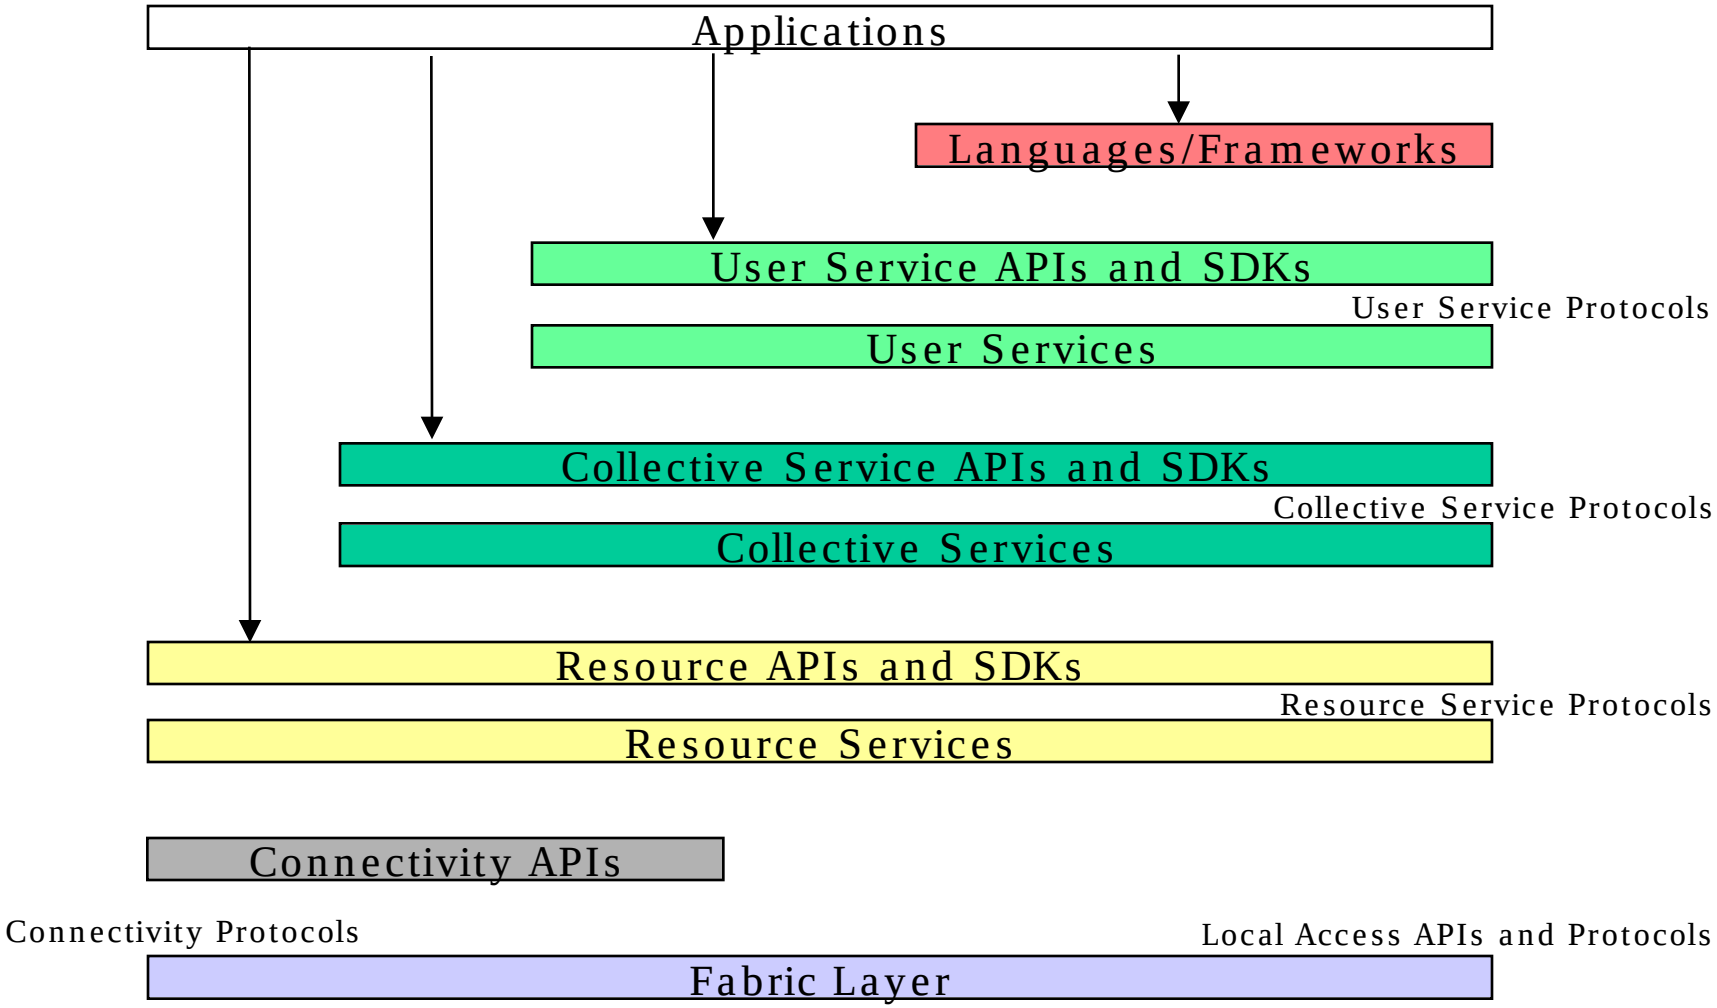

# Grid Services Architecture: Fabric Layer Protocols & Services

- Just what you would expect: the diverse mix of resources that may be shared
  - Individual computers, Condor pools, file systems, archives, metadata catalogs, networks, sensors, etc., etc.
- Few constraints on low-level technology: connectivity and resource level protocols form the "neck in the hourglass"
- Defined by interfaces not physical characteristics

# Grid Services Architecture: Connectivity Layer Protocols & Services

- Communication
  - Internet protocols: IP, DNS, routing, etc.
- Security: Grid Security Infrastructure (GSI)
  - Uniform authentication & authorization mechanisms in multi-institutional setting
  - Single sign-on, delegation, identity mapping
  - Public key technology, SSL, X.509, GSS-API
  - Supporting infrastructure: Certificate Authorities, key management, etc.

---

GSI: www.globus.org

# Grid Services Architecture: Resource Layer Protocols & Services

- Grid Resource Allocation Mgmt (GRAM)
  - Remote allocation, reservation, monitoring, control of compute resources
- GridFTP protocol (FTP extensions)
  - High-performance data access & transport
- Grid Resource Information Service (GRIS)
  - Access to structure & state information
- Network reservation, monitoring, control
- All integrated with GSI: authentication, authorization, policy, delegation

GRAM, GridFTP, GRIS: www.globus.org

# Grid Services Architecture: Collective Layer Protocols & Services

- Index servers aka metadirectory services
  - Custom views on dynamic resource collections assembled by a community
- Resource brokers (e.g., Condor Matchmaker)
  - Resource discovery and allocation
- Replica catalogs
- Co-reservation and co-allocation services
- Etc., etc.

---

Metadirectory: www.globus.org; Condor: www.cs.wisc.edu/condor

# Example: User Portal

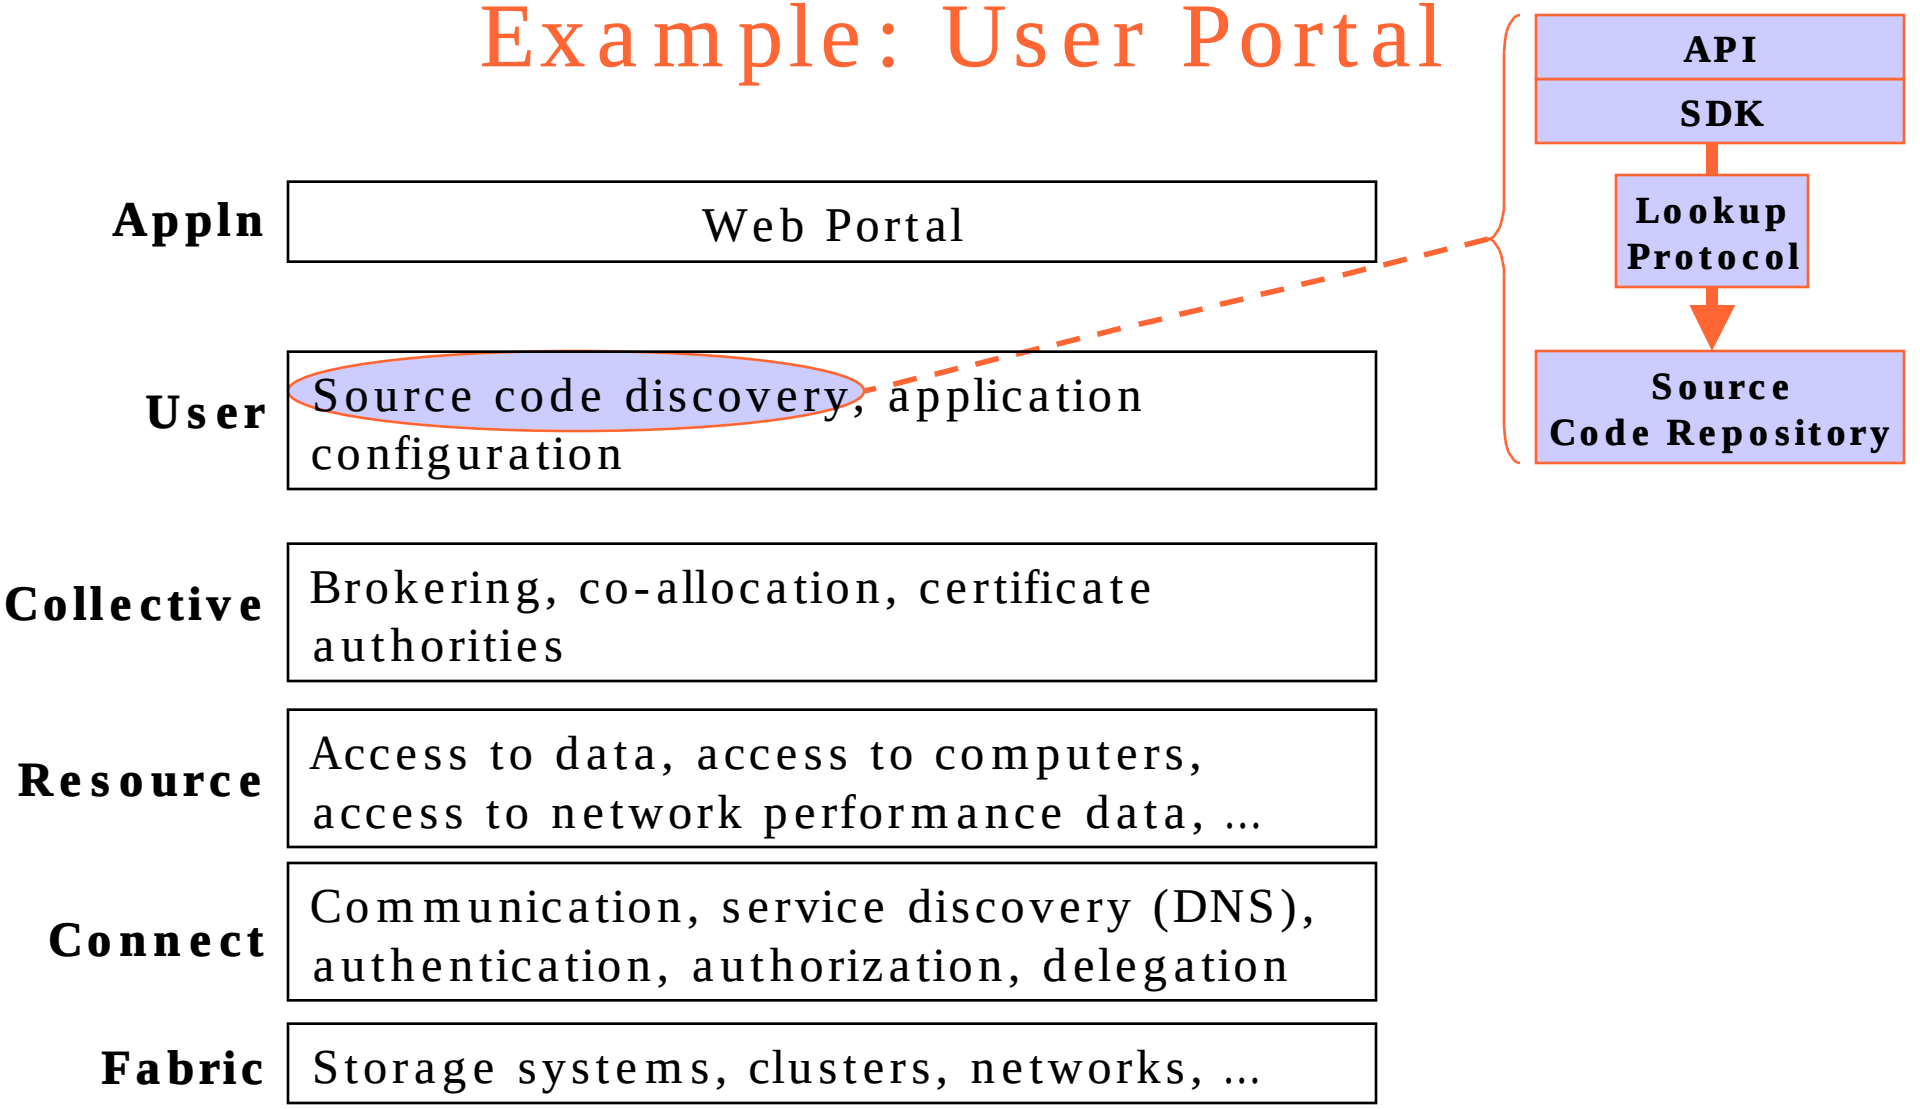

# Example: User Portal

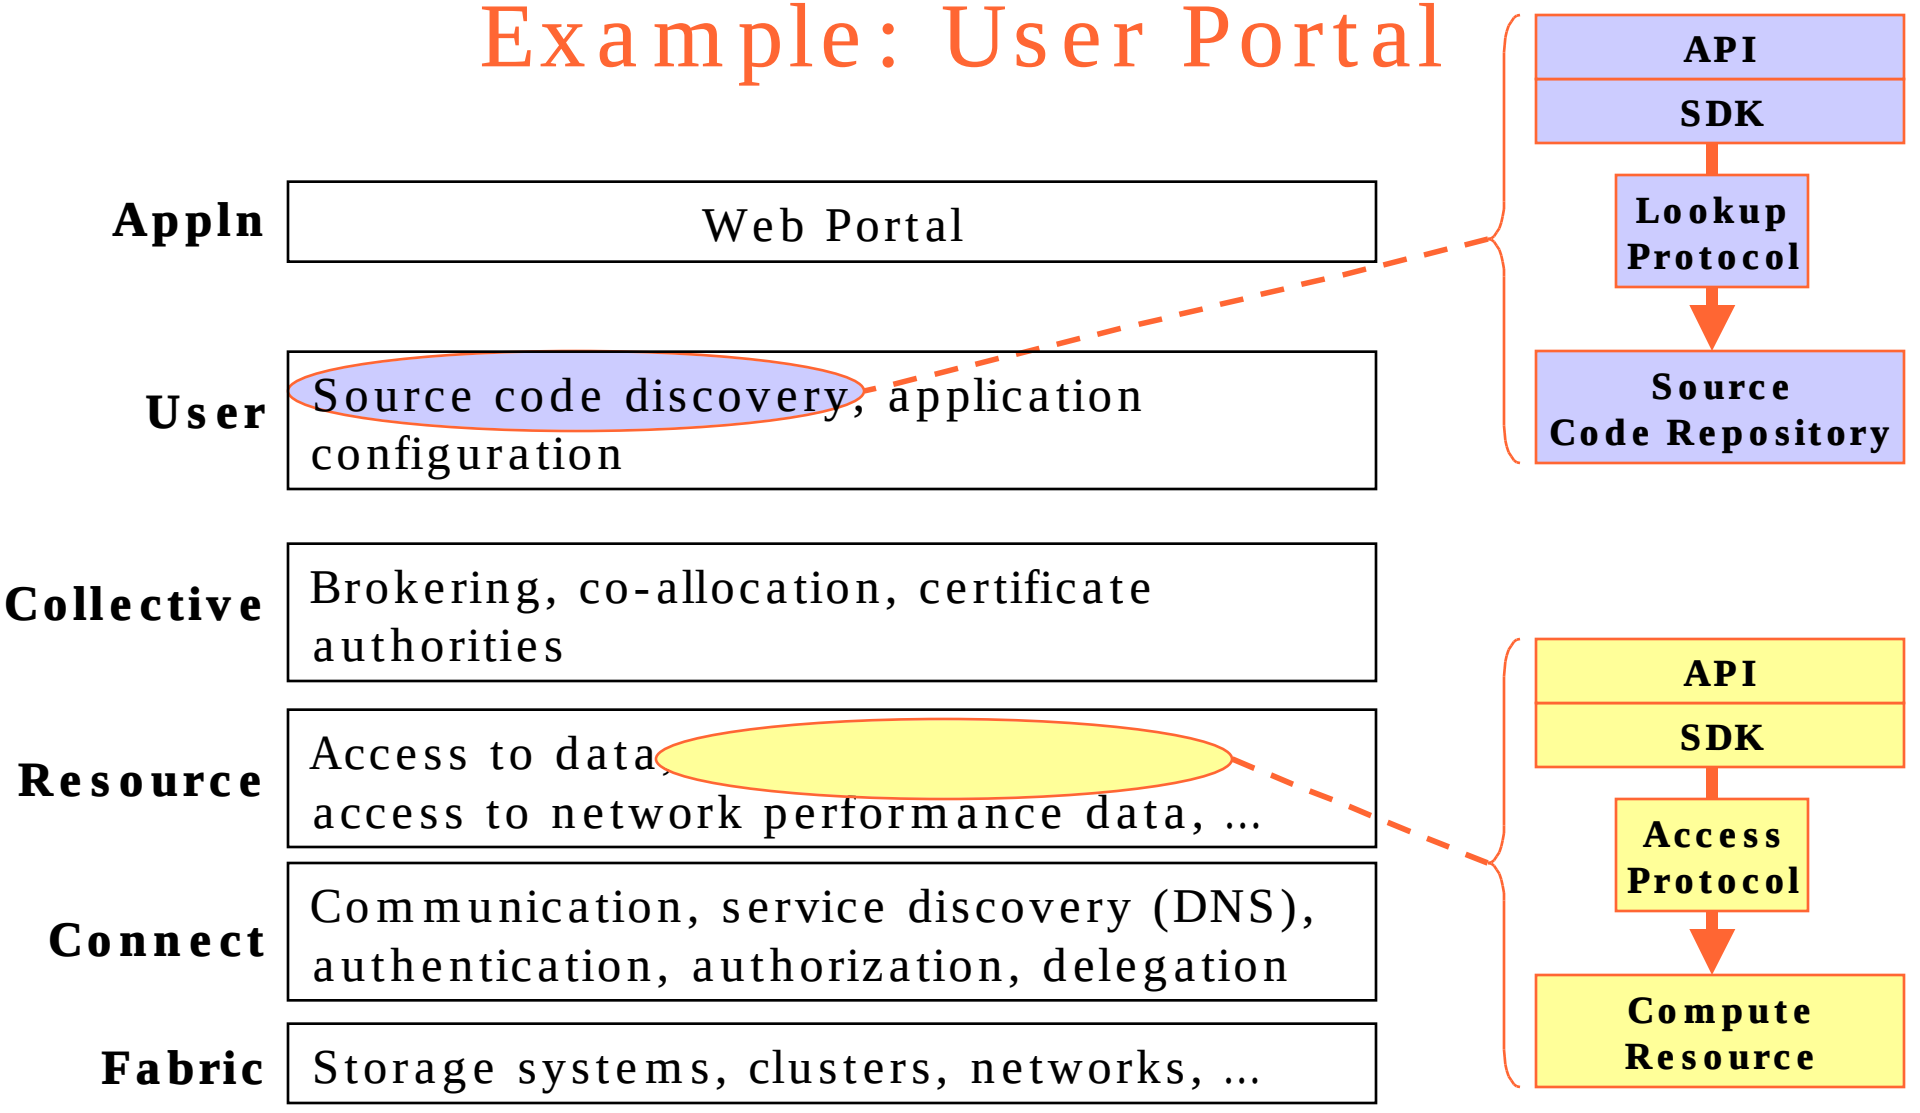

# Example: High-Throughput Computing System

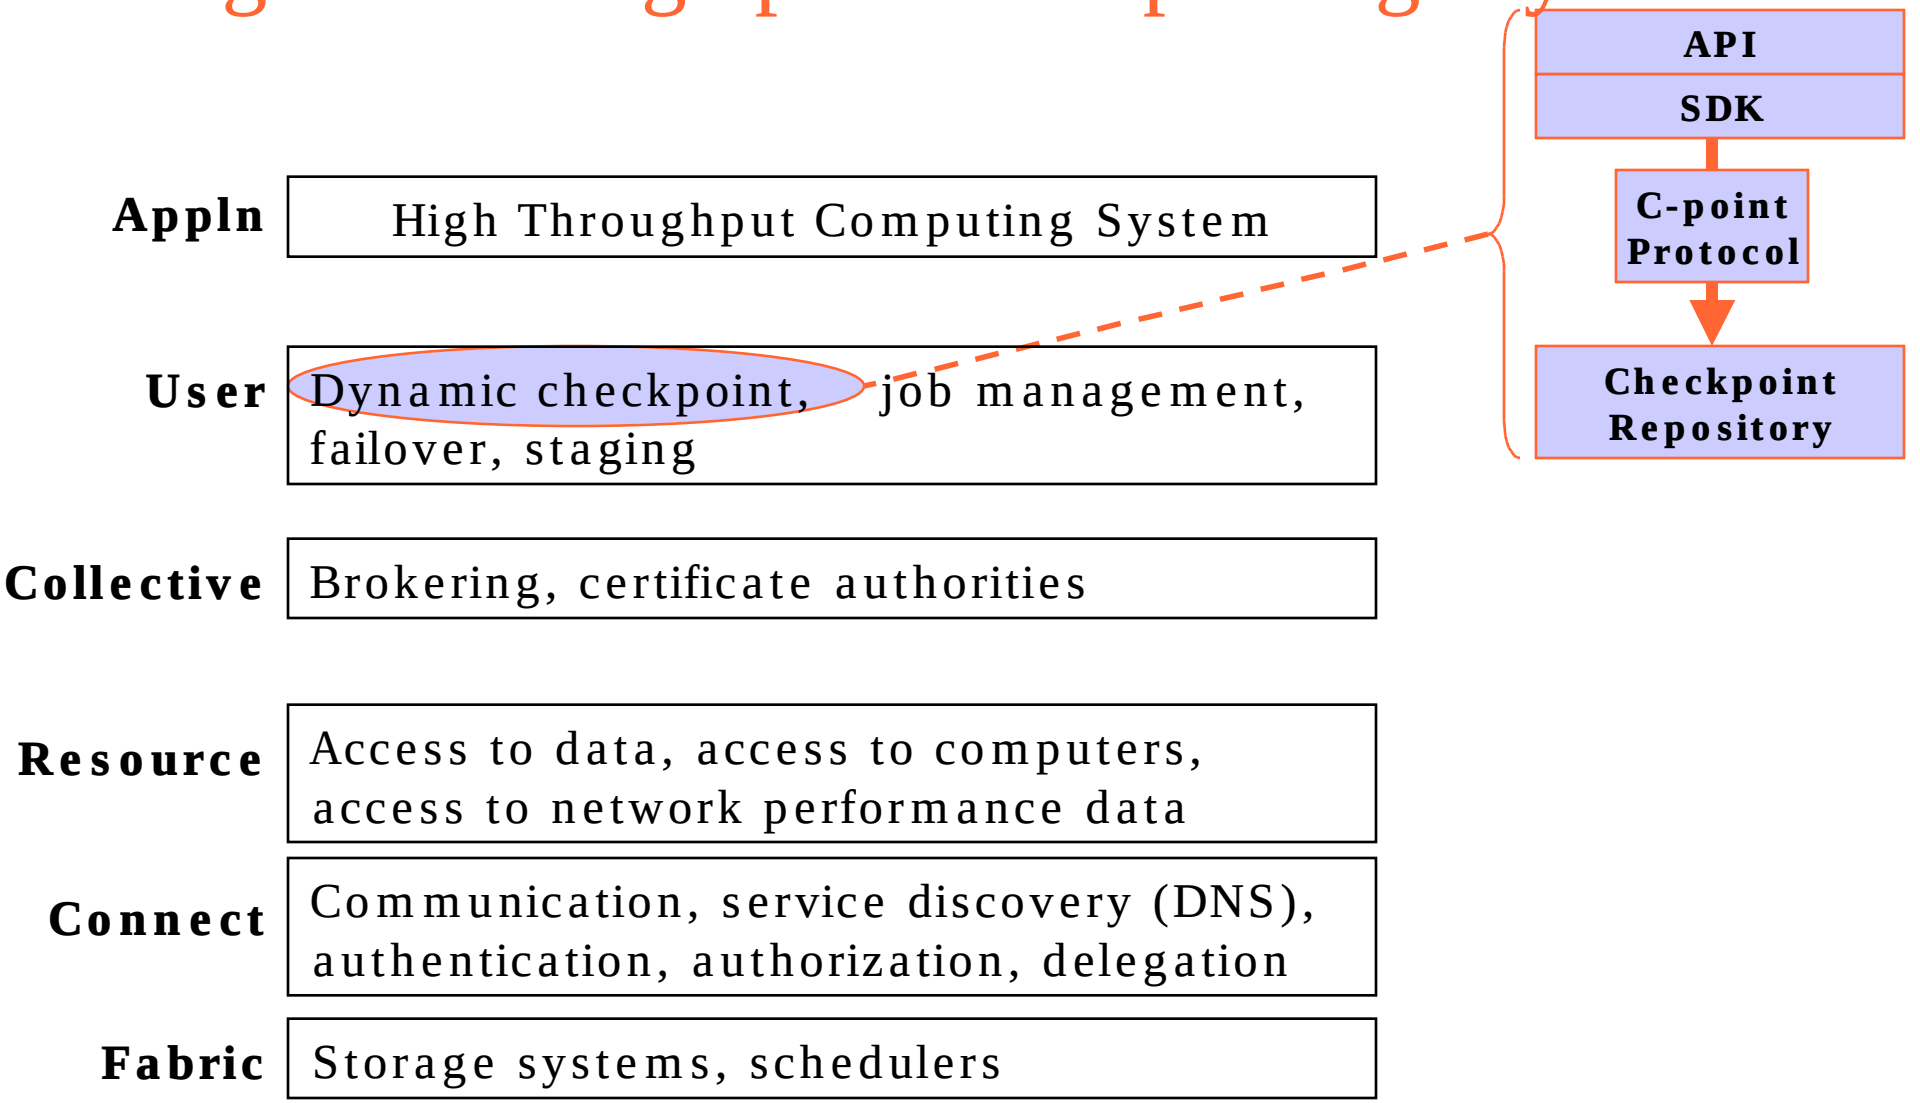

# Example: High-Throughput Computing System

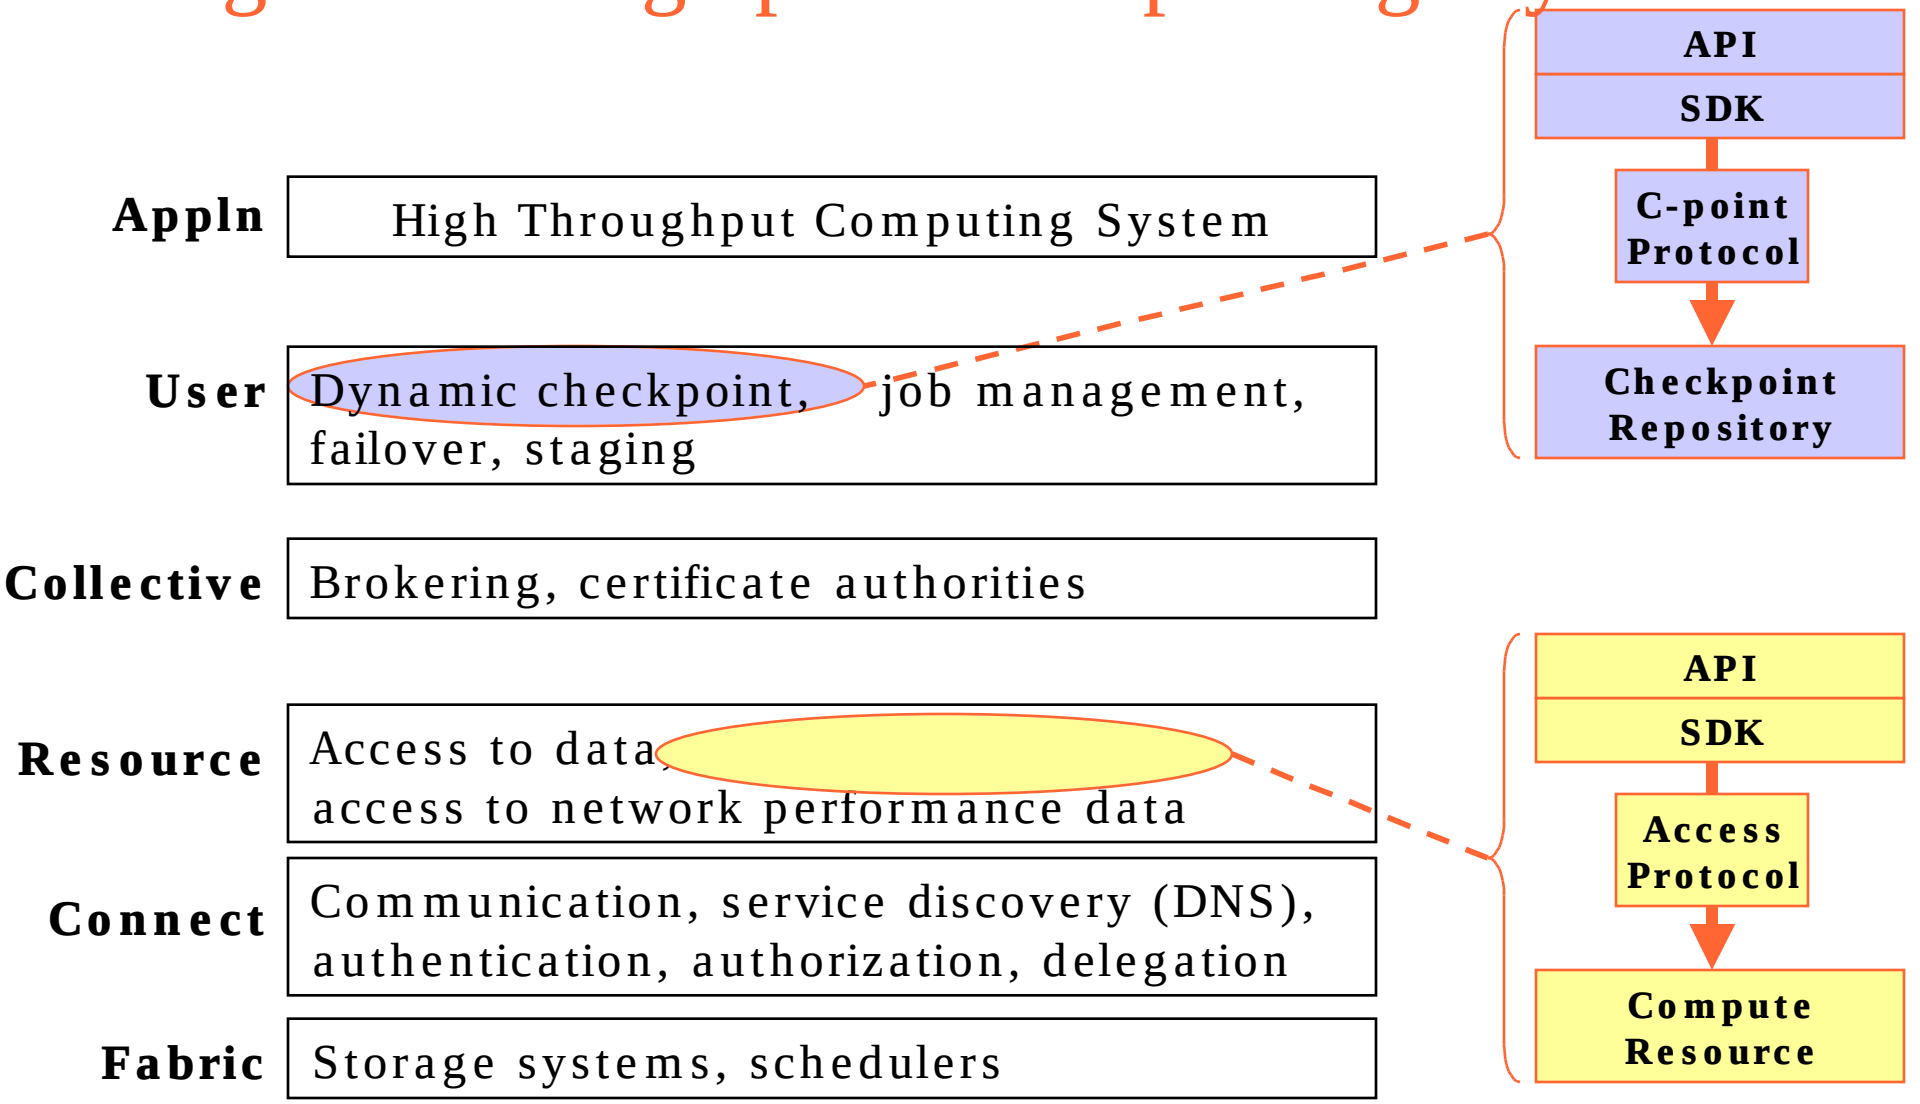

# Data Grids

# Data Grid Problem

- "Enable a geographically distributed community [of thousands] to pool their resources in order to perform sophisticated, computationally intensive analyses on Petabytes of data"
- Note that this problem
  - Is not unique to physics
  - Overlaps strongly with other Grid problems
- Data Grids do introduce new requirements and R&D challenges

# Data Grid Hierarchy Concept

~PBytes/sec

Online System

~100 MBytes/sec

Offline Processor Farm

~20 TIPS

~100 MBytes/sec

CERN Computer Centre

HPSS

~622 Mbits/sec
or Air Freight (deprecated)

France Regional Centre

Germany Regional Centre

Italy Regional Centre

FermiLab ~4 TIPS

~622 Mbits/sec

Tier2 Centre
~1 TIPS

*There is a "bunch crossing" every 25 nsecs.*

*There are 100 "triggers" per second*

*Each triggered event is ~1 MByte in size*

Physics data cache

Institute
~0.25TIPS

~1 MBytes/sec

Physicist workstations

*Physicists work on analysis "channels".*

*Each institute will have ~10 physicists working on one or more channels; data for these channels should be cached by the institute server*

# Major Data Grid Projects

- Earth System Grid (DOE Office of Science)
  - DG technologies, climate applications
- European Data Grid (EU)
  - DG technologies & deployment in EU
- GriPhyN (NSF ITR)
  - Investigation of “Virtual Data” concept
- Particle Physics Data Grid (DOE Science)
  - DG applications for HENP experiments

# Data Grid Architecture

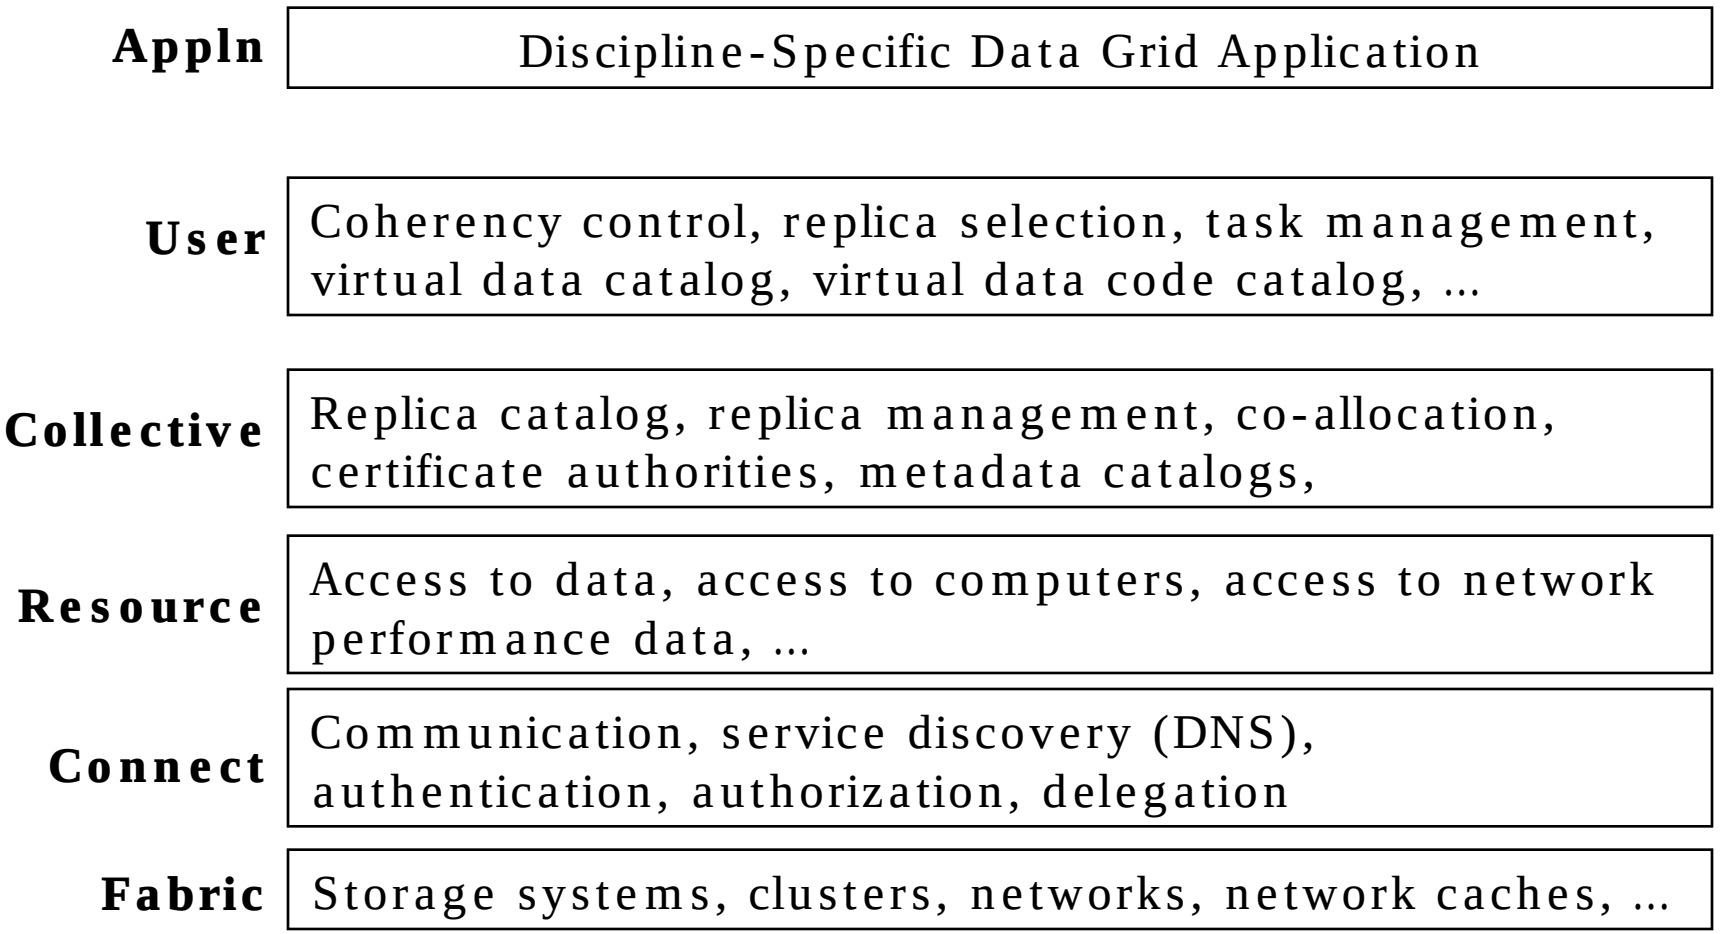

# Globus Data-Intensive Services Architecture

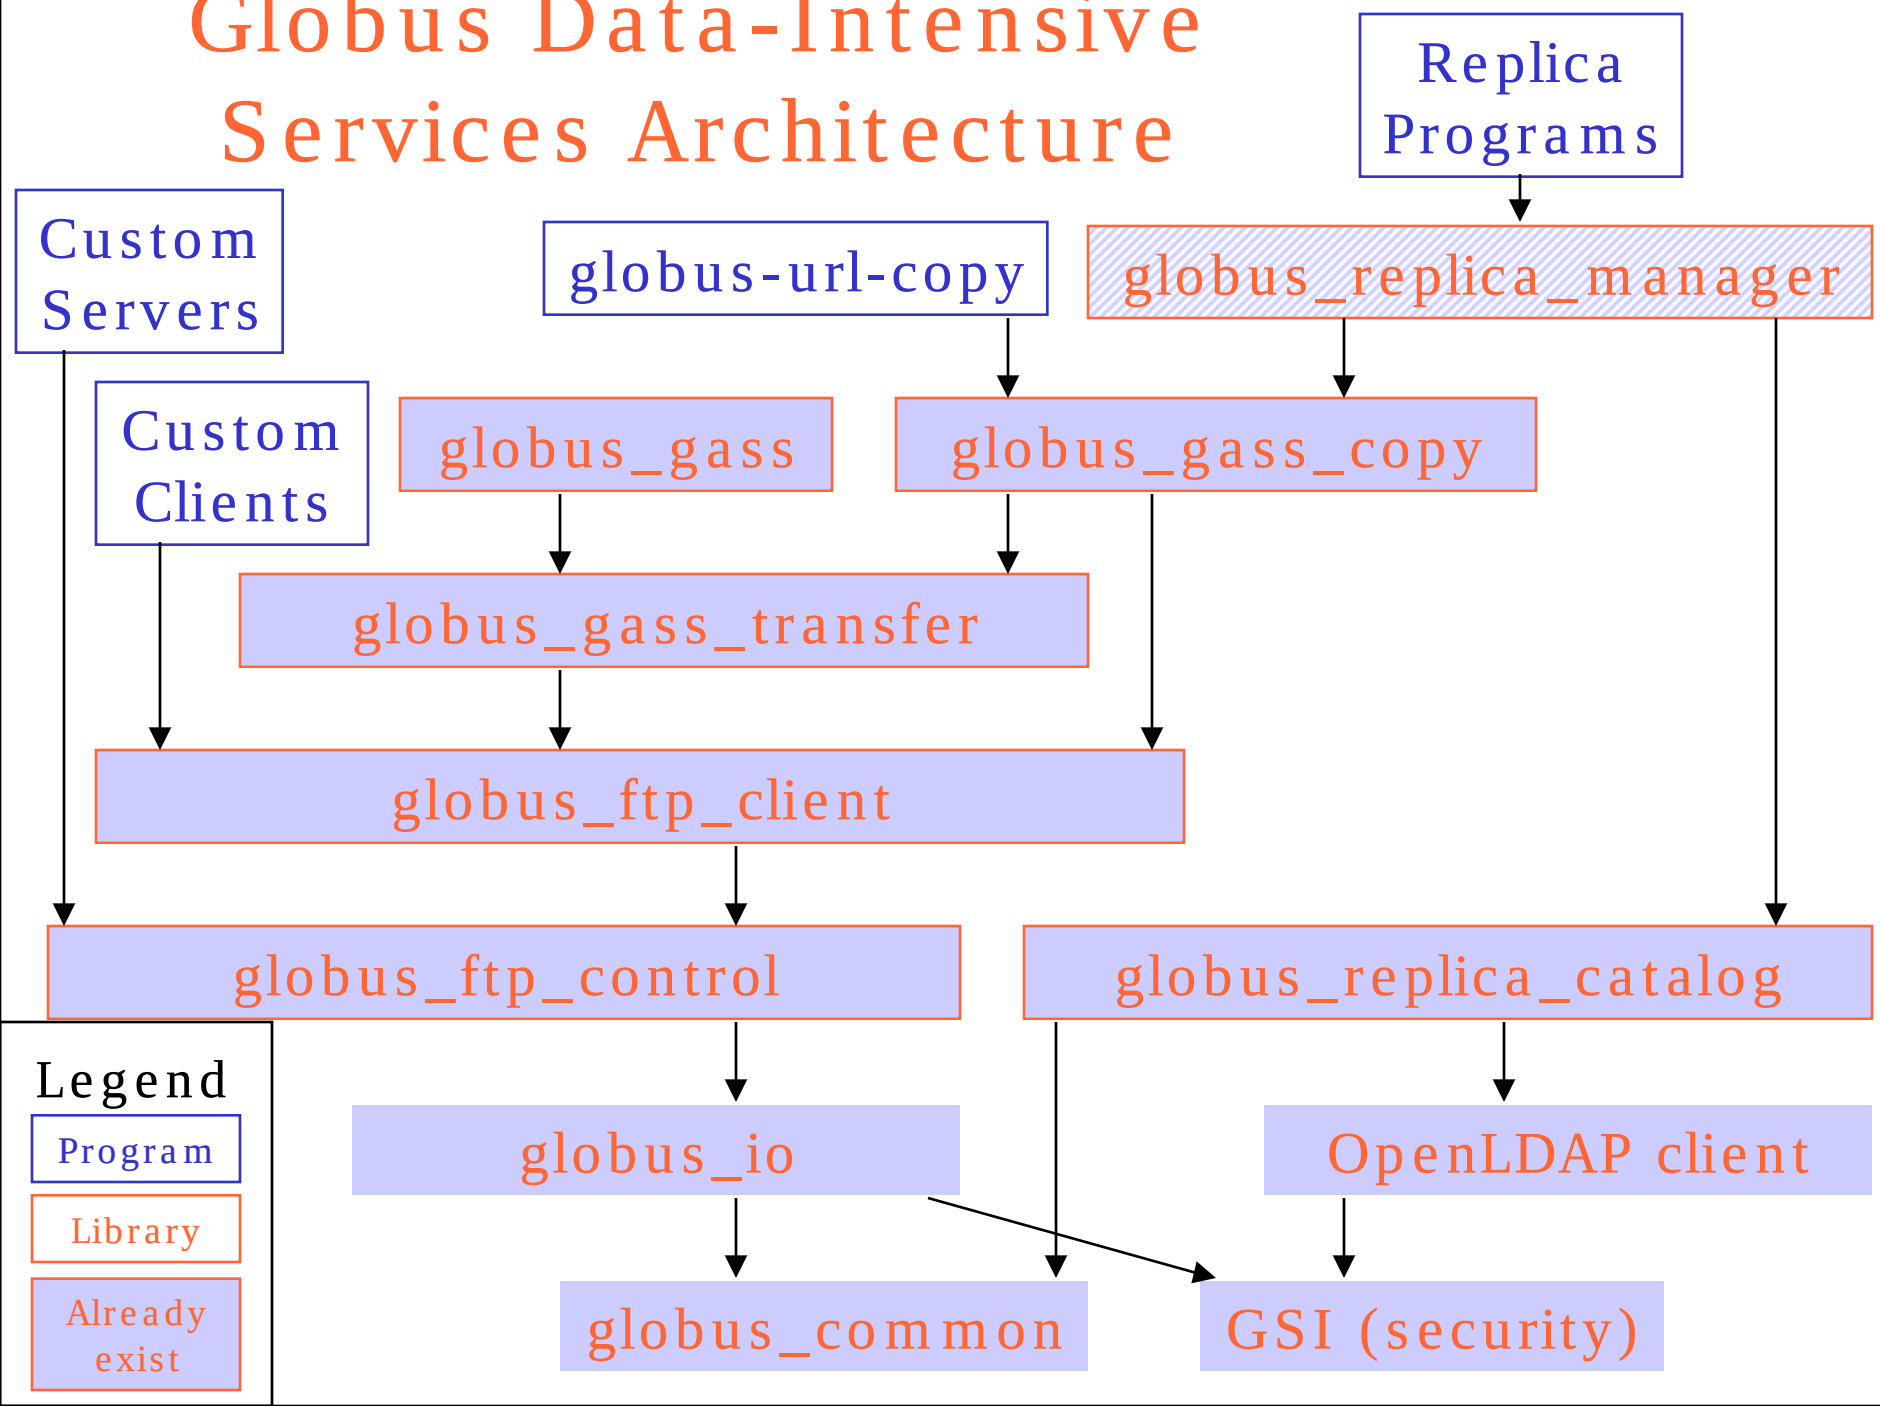

# High-Level View of Earth System Grid: A Model Architecture for Data Grids

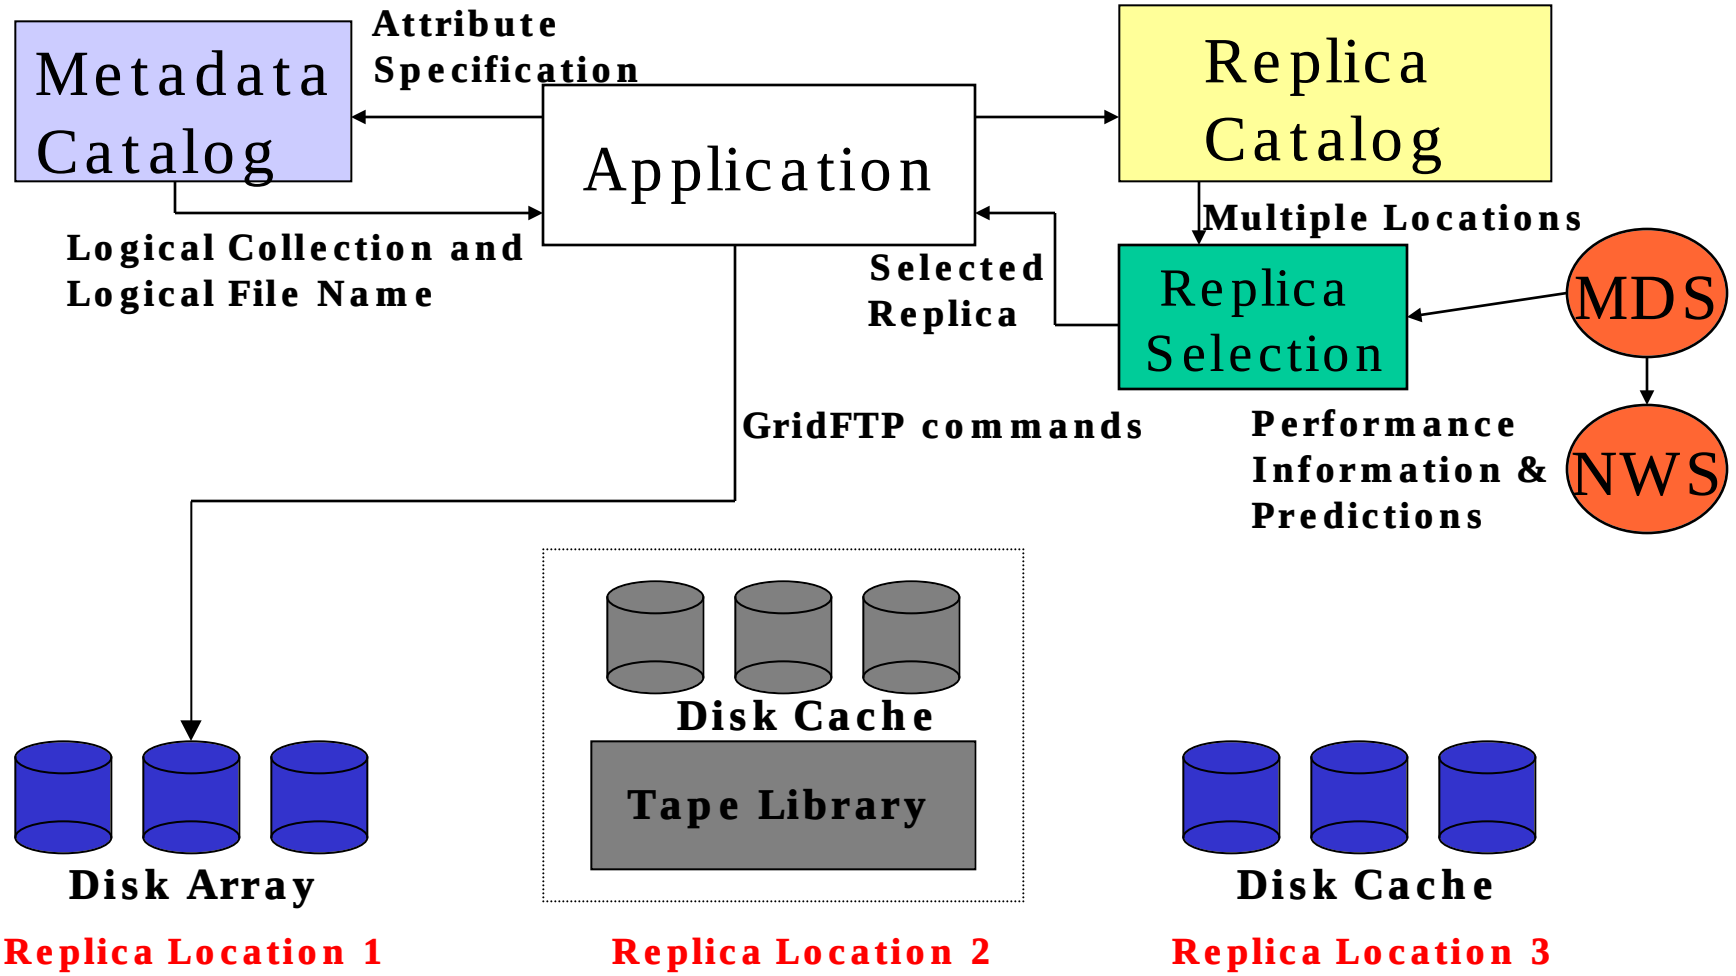

# Grid Physics Network (GriPhyN)

## Enabling R&D for advanced data grid systems, focusing in particular on Virtual Data concept

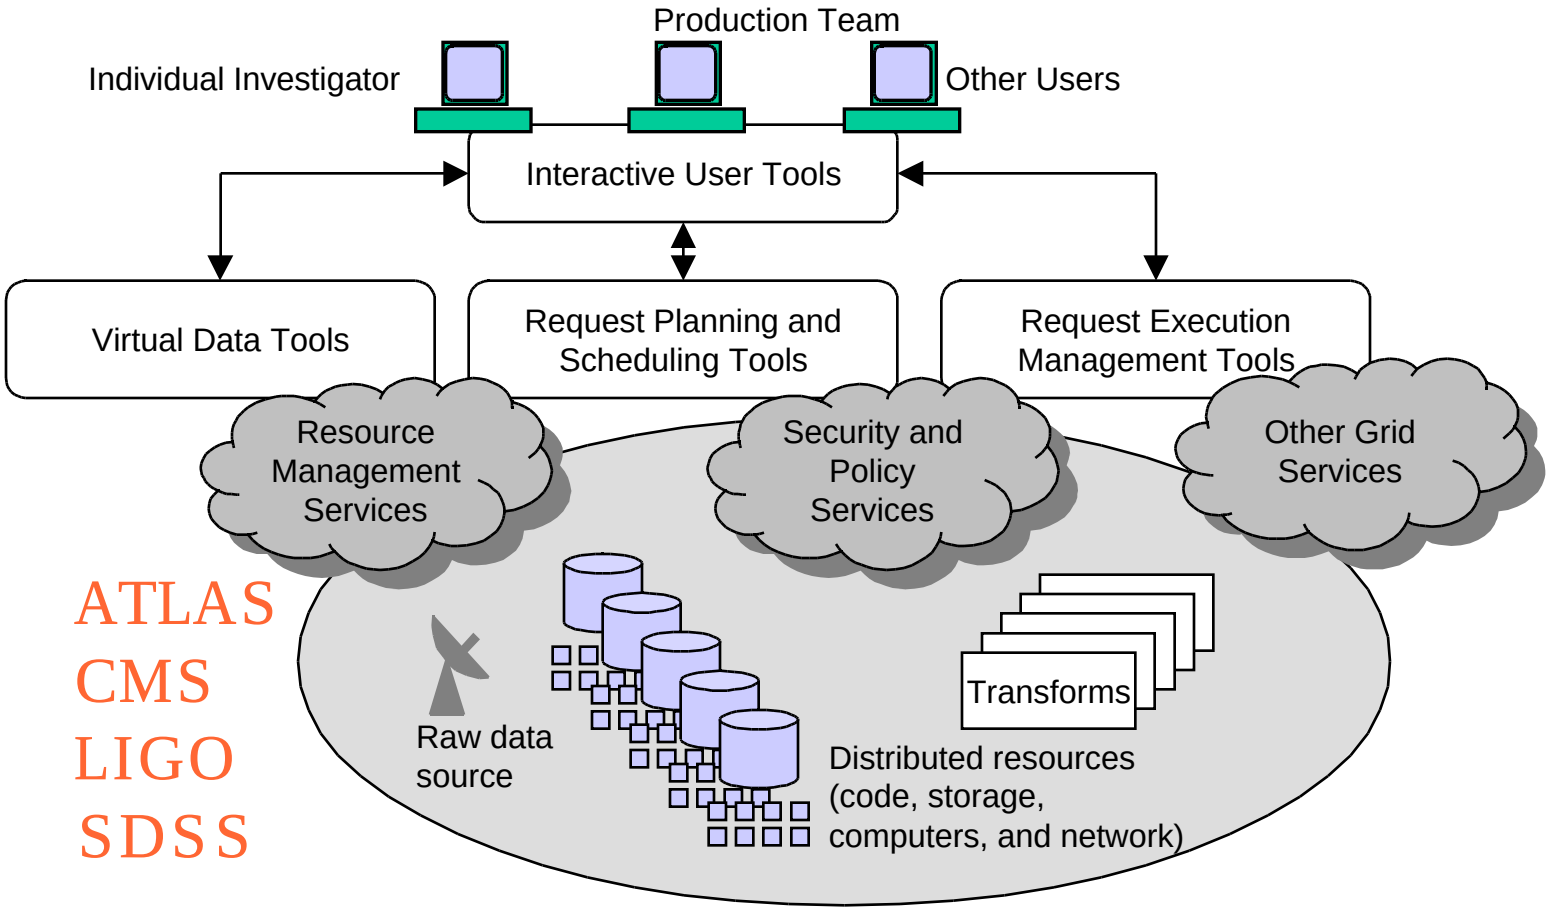


ATLAS
CMS
LIGO
SDSS

# The Virtual Data Concept

"[a virtual data grid enables] the definition and delivery of a potentially unlimited virtual space of data products derived from other data. In this virtual space, requests can be satisfied via direct retrieval of materialized products and/or computation, with local and global resource management, policy, and security constraints determining the strategy used."

# Virtual Data in Action

- Data request may
  - Access local data
  - Compute locally
  - Compute remotely
  - Access remote data
- Scheduling subject to local & global policies
- Local autonomy

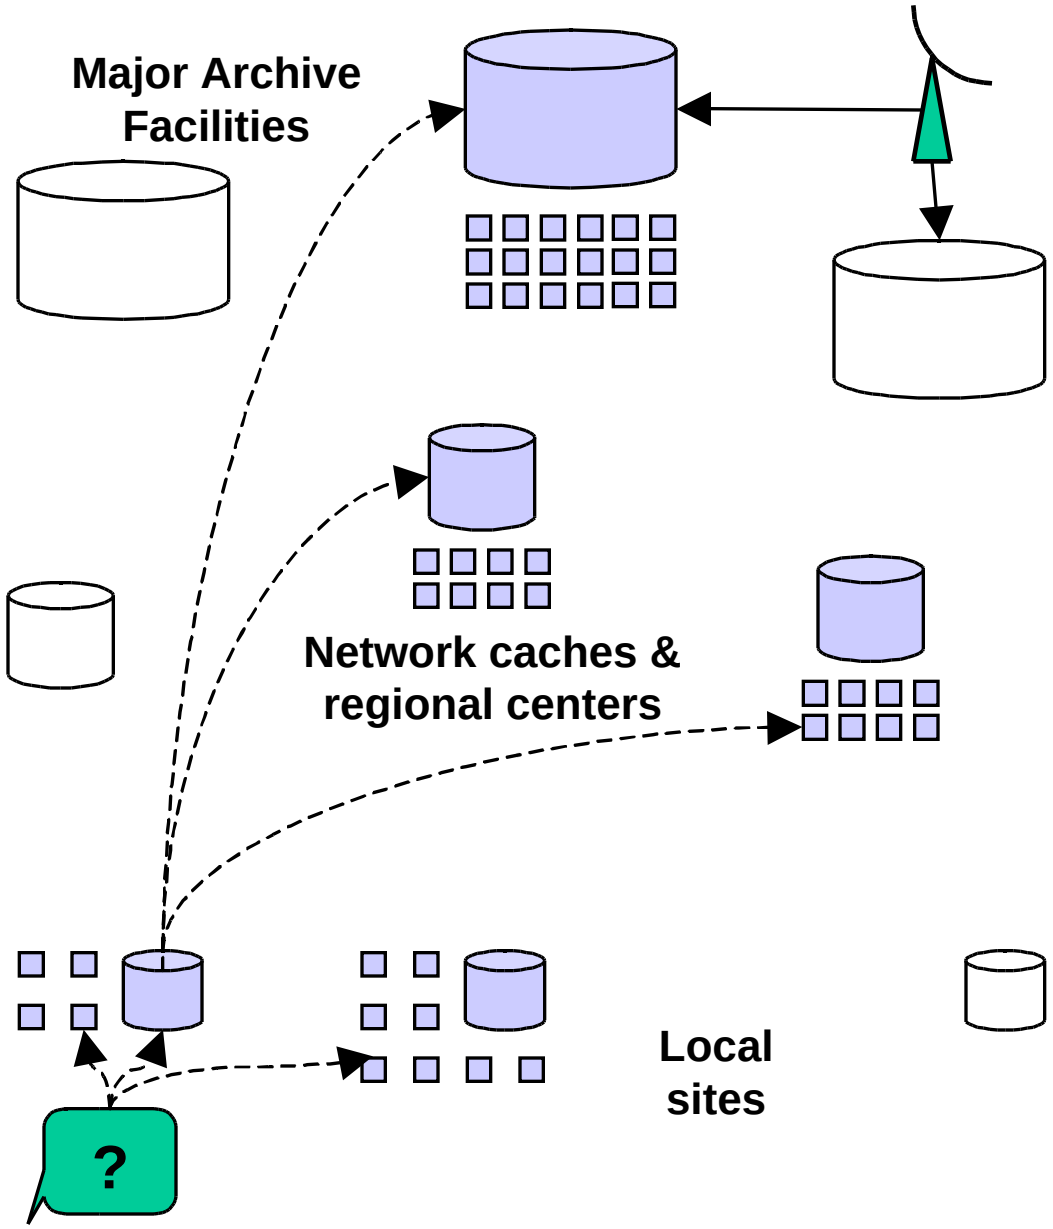

# Abbreviated Acknowledgments

- Other Globus principals include: Steve Tuecke, Ann Chervenak, Bill Allcock, Gregor von Laszewski, Lee Liming, Steve Fitzgerald
- Data Grid collaborators: Arie Shoshani, Reagan Moore, Miron Livny
- GriPhyN Co-PI: Paul Avery
- PPDG PIs: Richard Mount, Harvey Newman

# Summary

- Grid computing concepts are real: they underlie a variety of major projects in physics, other sciences, and industry
- Promise is that by allowing coordinating sharing, enable quantitative & qualitative changes in problem solving techniques
- Clear picture of Grid architecture emerging
- Data Grid concepts and technologies are less mature but are appearing rapidly

# For More Information

- Book (Morgan Kaufman)
  - www.mkp.com/grids
- Globus
  - www.globus.org
- Grid Forum
  - www.gridforum.org
- PPDG
  - www.ppdg.net
- GriPhyN
  - www.griphyn.org

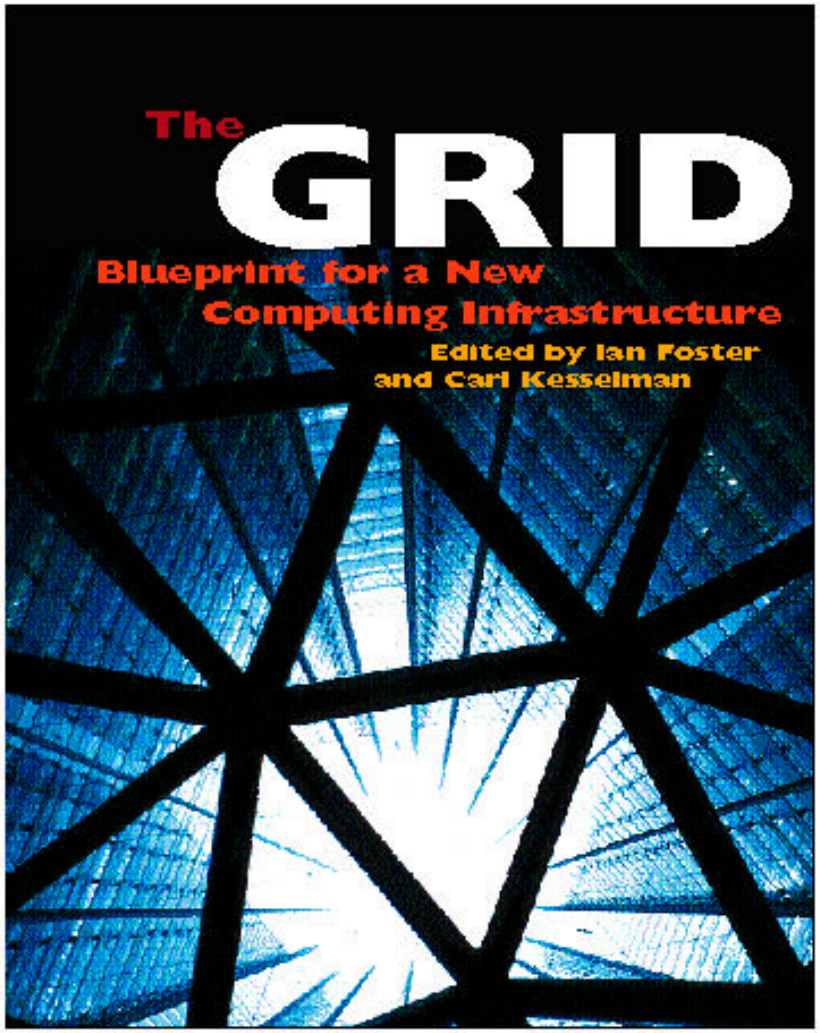

The End