Grid Computing Research

Ian Foster

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

Overview

- Grid computing & its importance to DOE
- Grid R&D at Argonne
- Defining and instantiating a Grid Services architecture
- Data Grids
- Future directions

The Grid: The Web on Steroids

Web: Uniform access to HTML documents

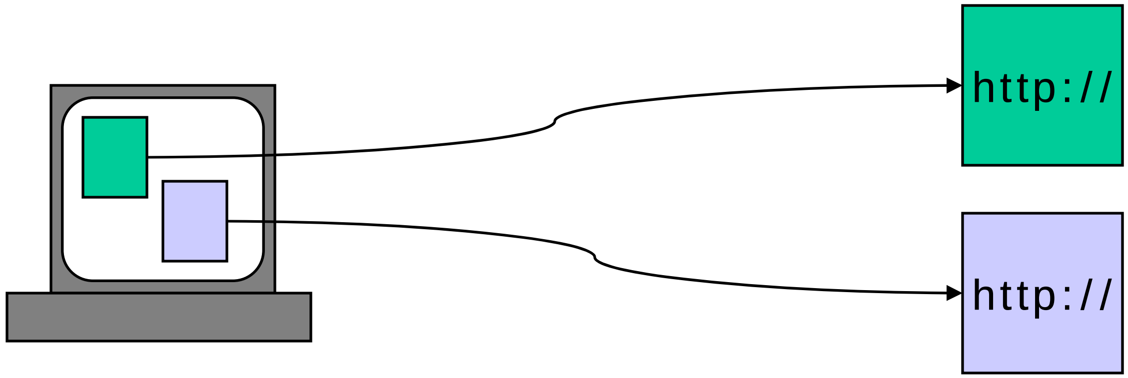

Grid: Flexible, high-perf access to all significant resources

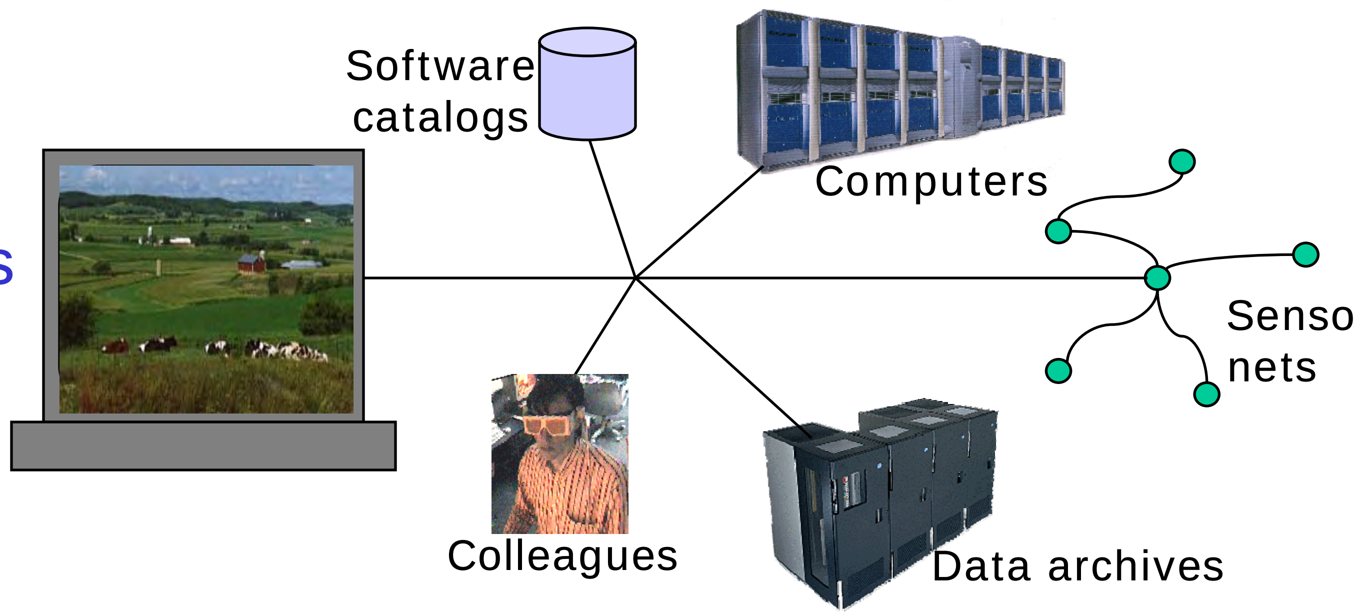

On-demand creation of powerful virtual computing systems

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

Grids and DOE

- Distinctive characteristics of much DOE science encourages Grid concepts
  - Unique and expensive facilities: accelerators, microscopes, supercomputers, ...
  - Large-scale, multidisciplinary science: climate, materials, high energy physics, ...
  - Large-scale simulation, data-intensive science
- The question is not *whether* to Grid-enable DOE Science—but *how*

Grid Communities & Applications: Data Grids for High Energy Physics

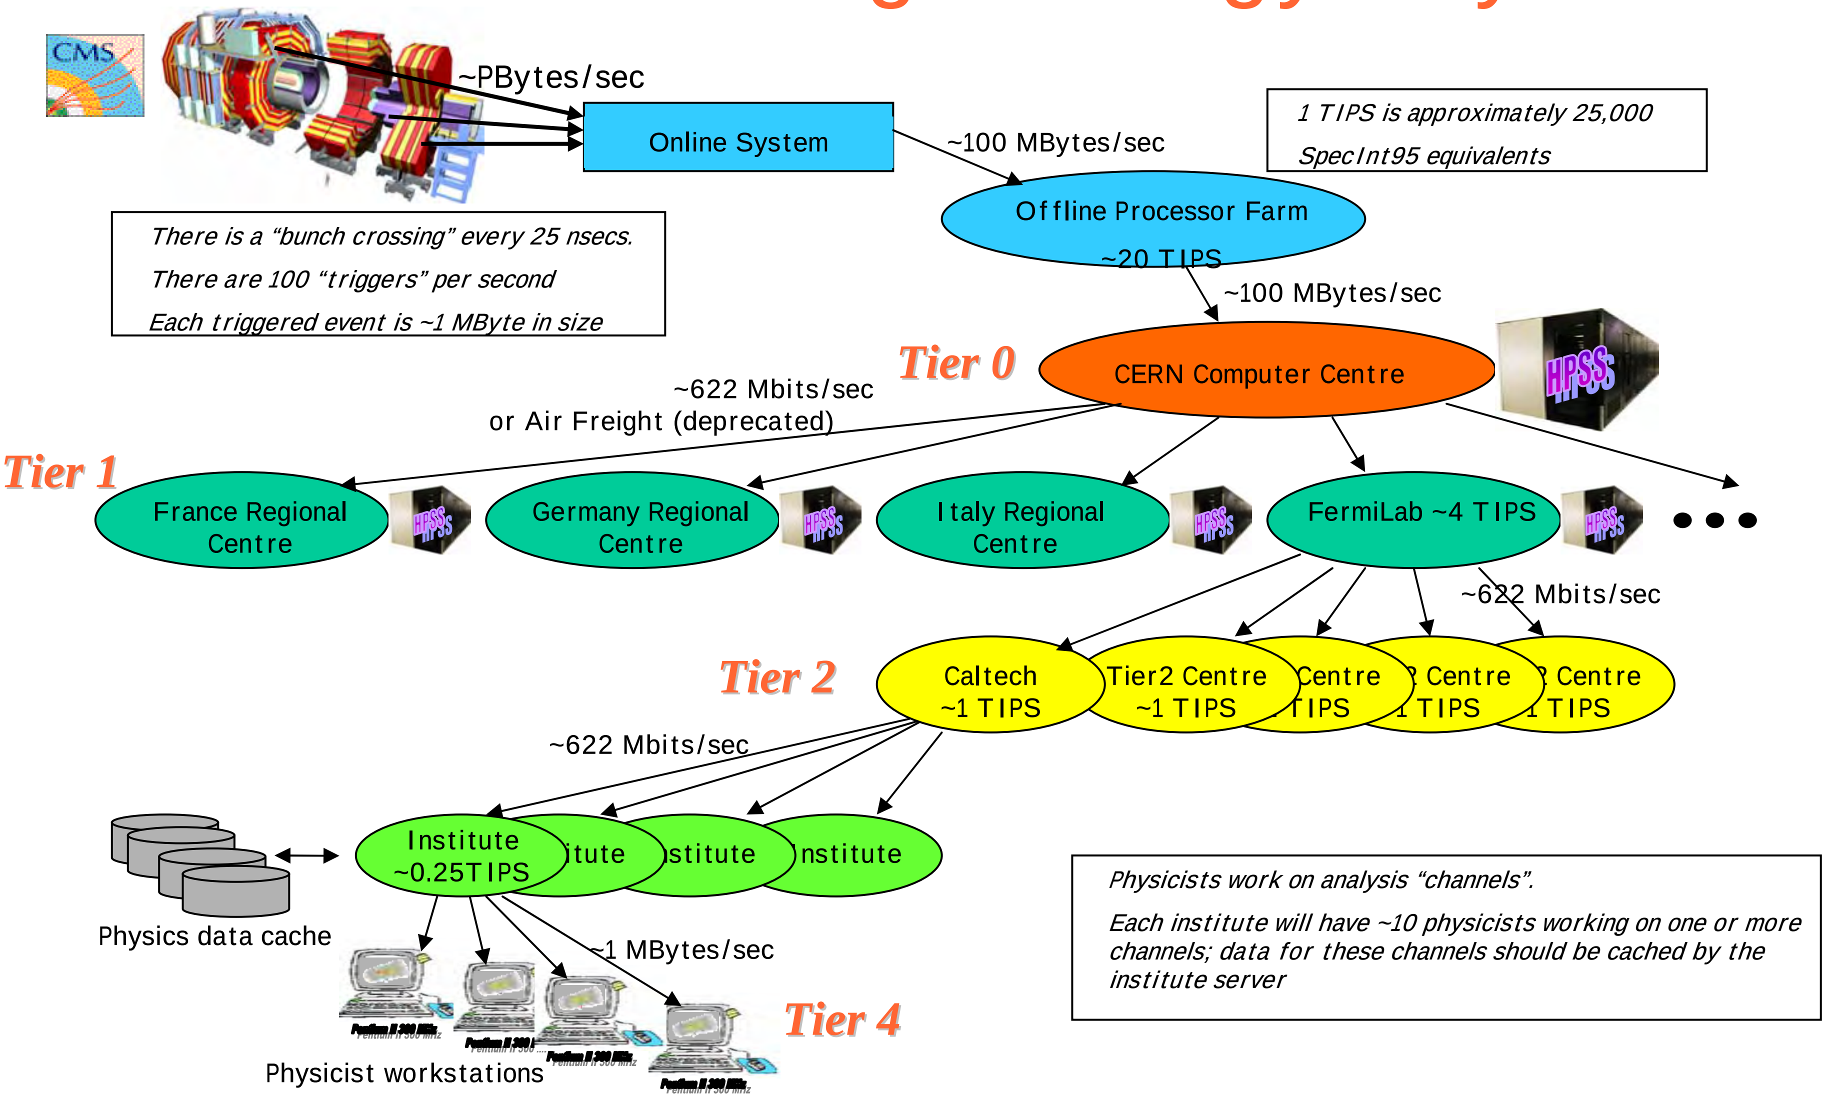

Image courtesy Harvey Newman, Caltech

Grid Communities & Applications: Online Instrumentation

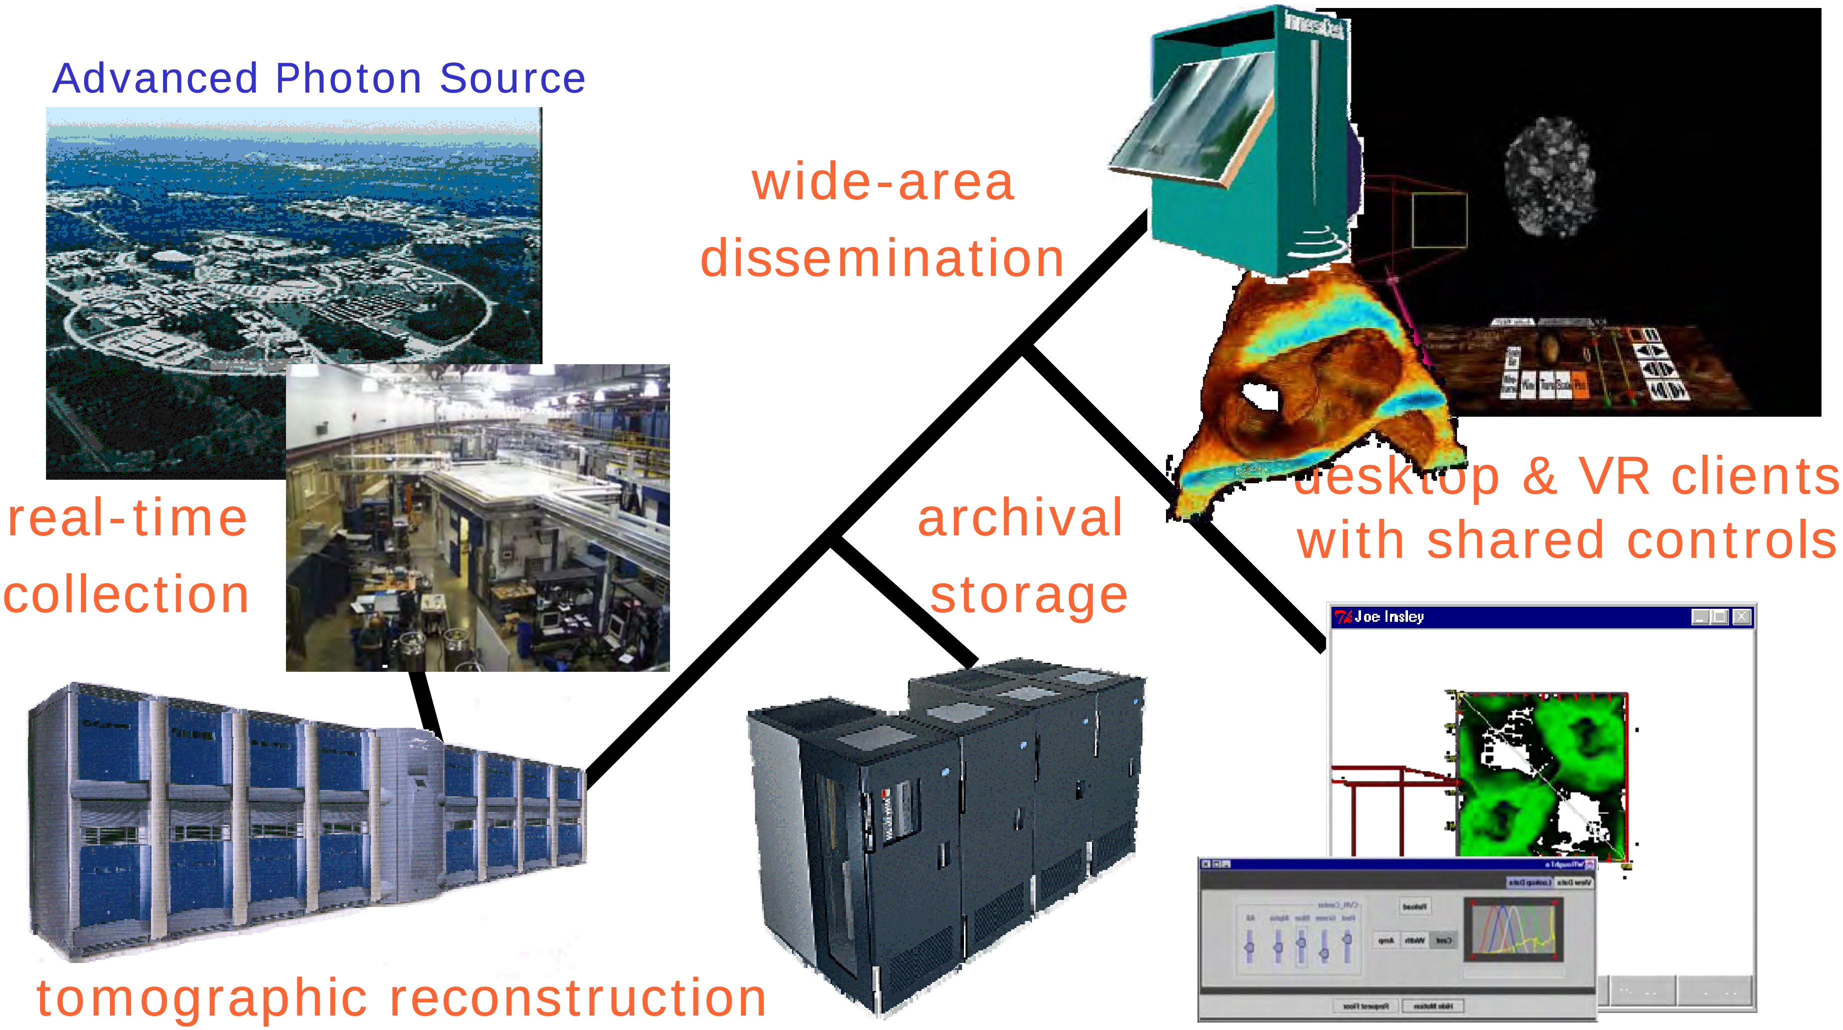

DOE X-ray grand challenge: ANL, USC/ISI, NIST, U.Chicago

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

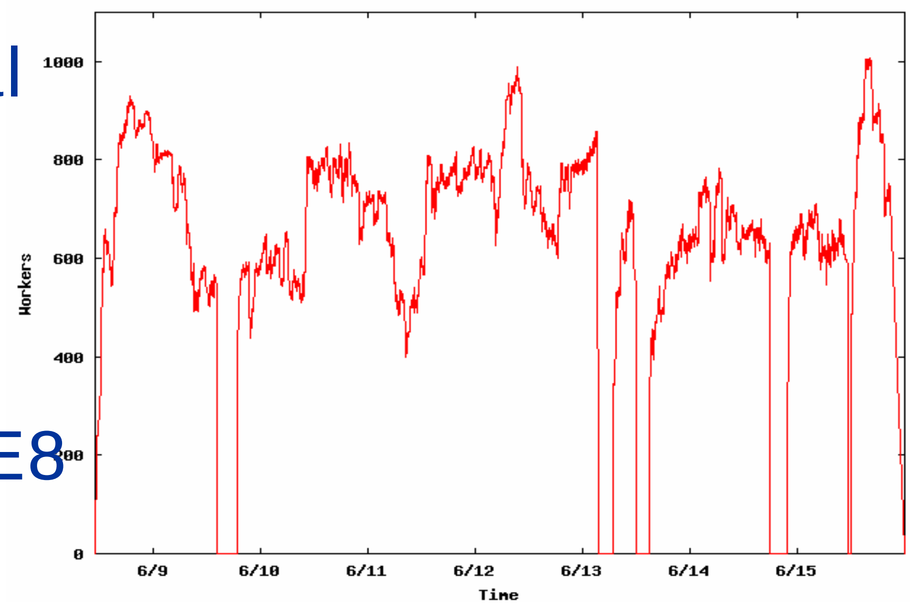

14,5,28,24,1,3,16,15,
10,9,21,2,4,29,25,22,
13,26,17,30,6,20,19,
8,18,7,27,12,11,23

MetaNEOS: Argonne, Iowa, Northwestern, Wisconsin

Grid R&D at Argonne

- Started in 1995 (I-WAY experiment, s/w)
- Globus R&D (much joint with USC/ISI)
  - Innovative security, resource management, data access, information, communication, fault detection, etc., etc. technologies
  - Large user base among tool developers
  - Widespread adoption in “production” grids: e.g., NASA IPG, NSF NTG, DOE DISCOM
  - Exciting application demonstrations
- Access Grid collaboration technologies

Access Grid Collaboration Technology

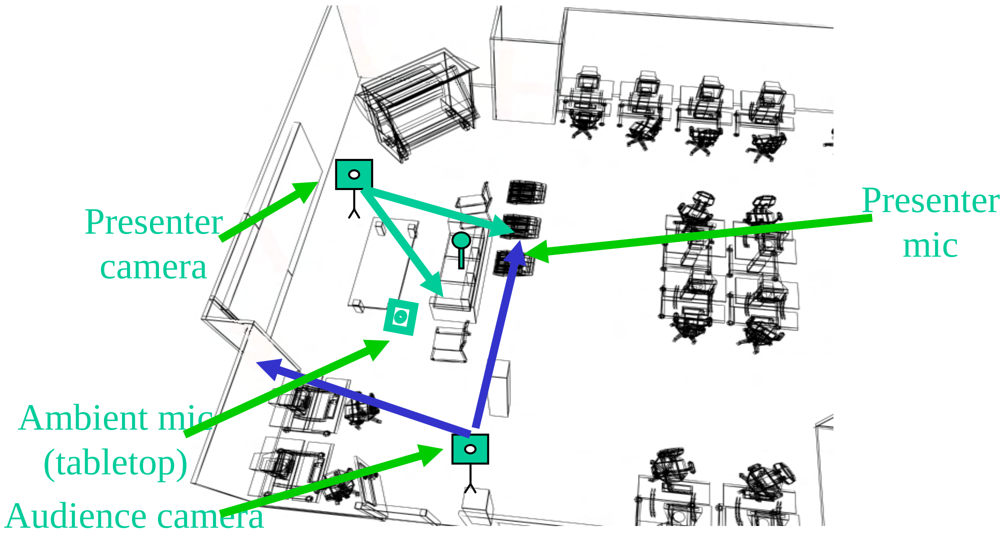

- Designed spaces for group interactions
- Hands free audio
- Multiple video and audio streams
- Wide field of view

Addressing Cross-Cutting Technical Issues

- 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)

“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

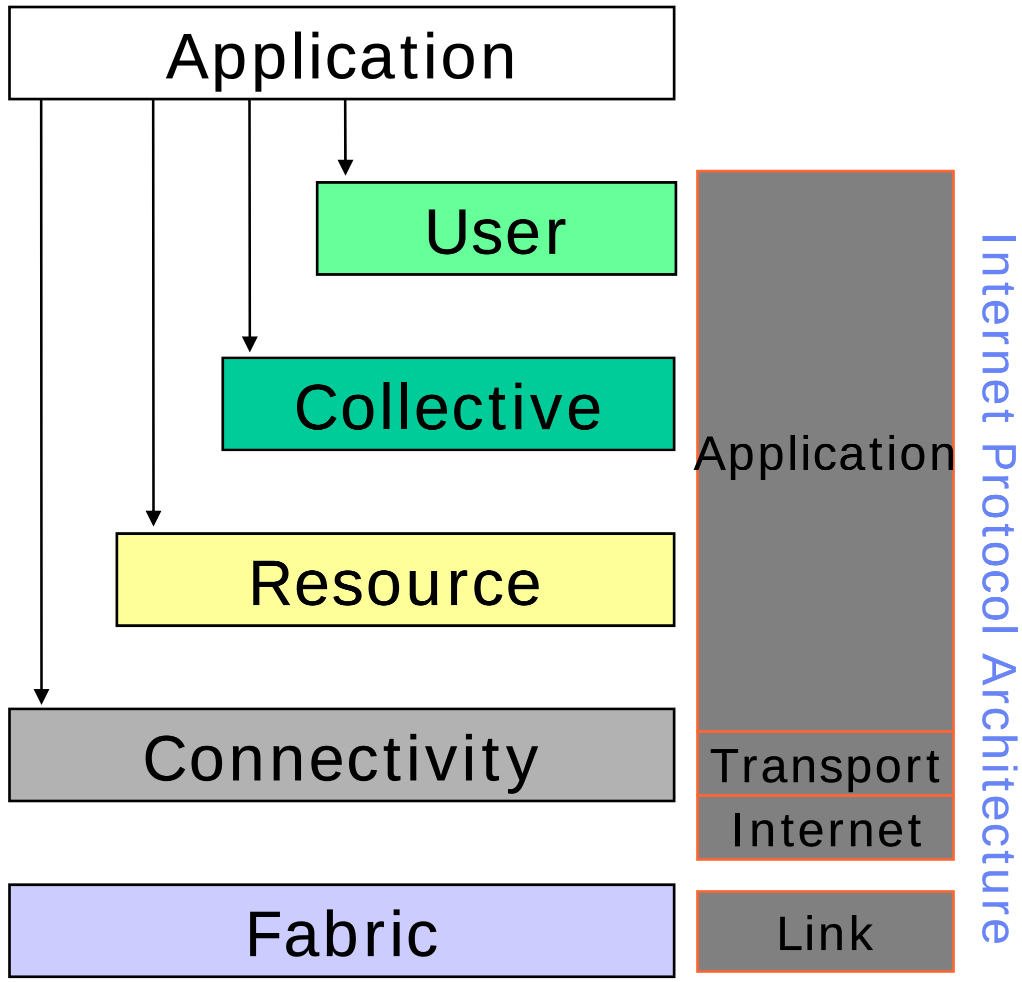

Instantiating the Grid Architecture: Connectivity Layer Protocols & Services

- 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

Instantiating the Grid 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

Instantiating the Grid 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

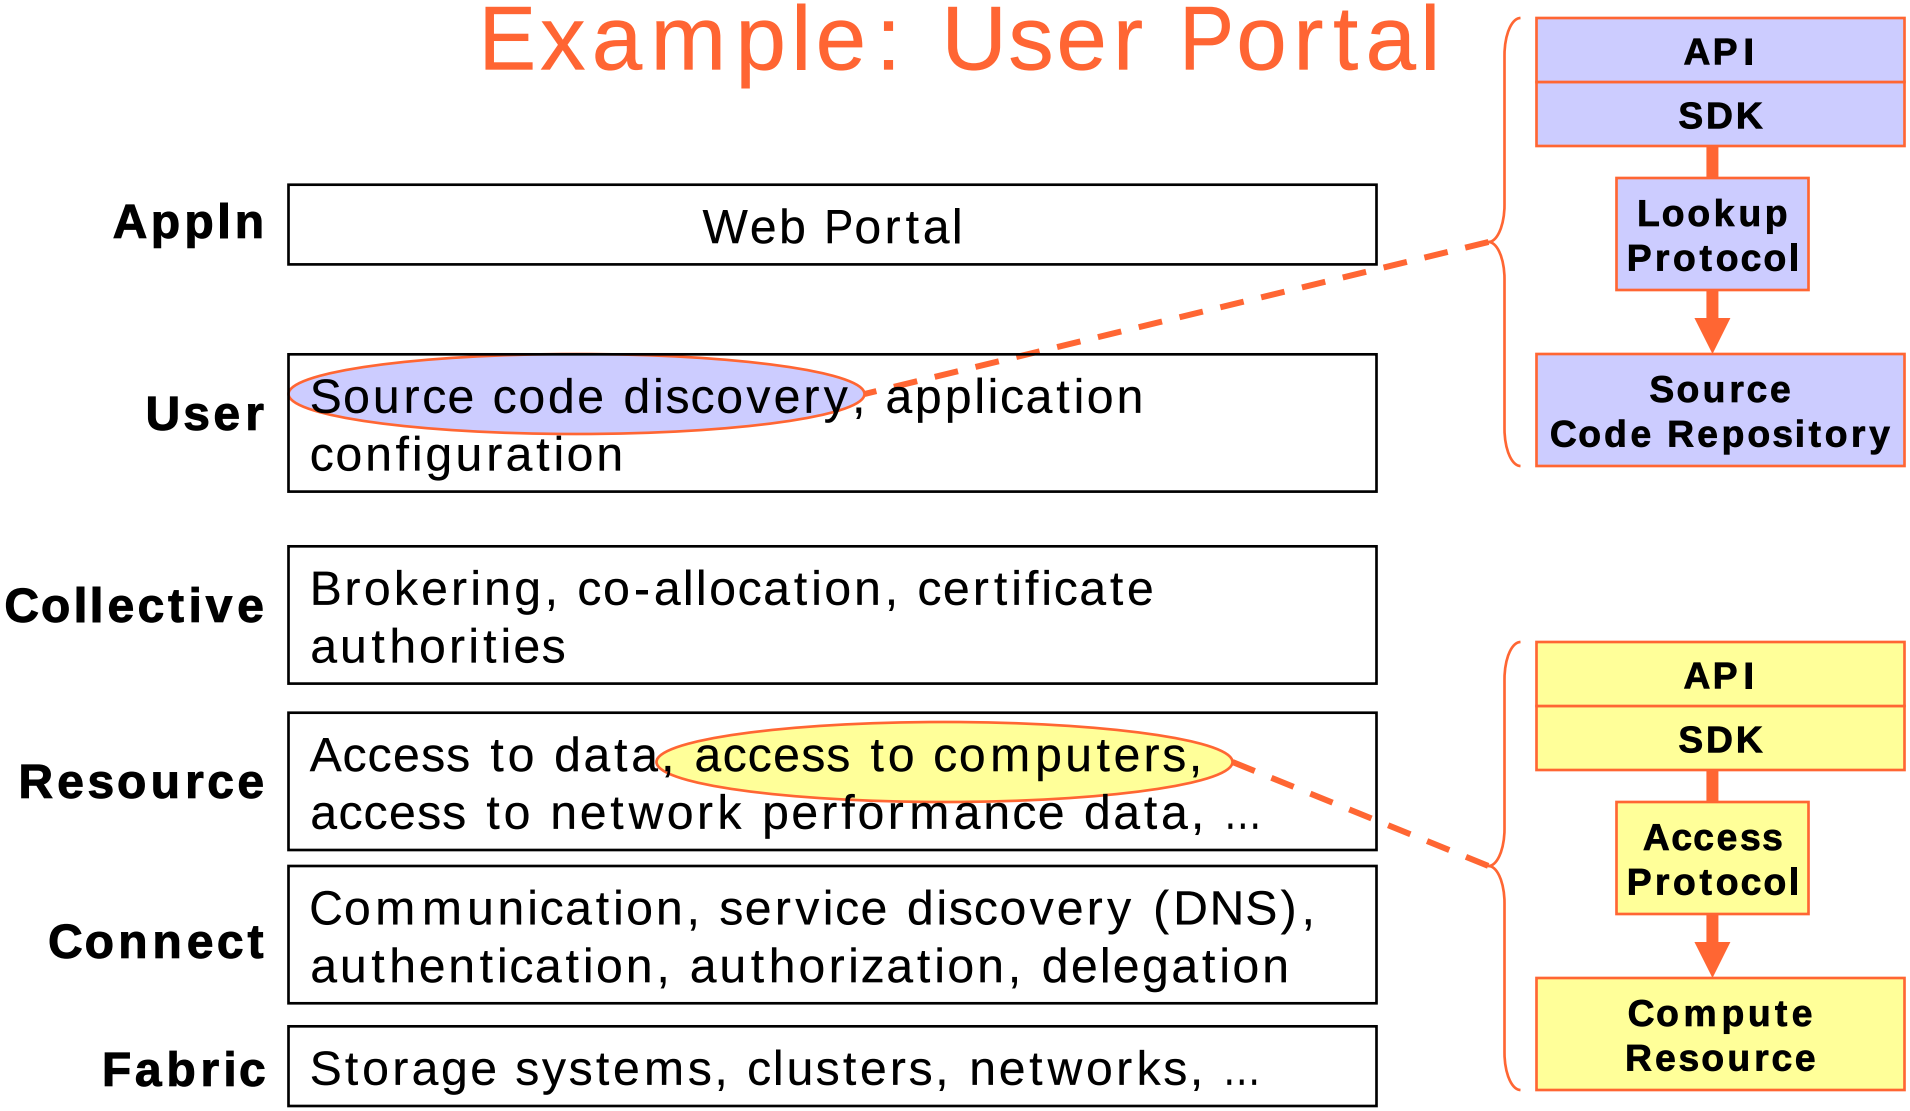

Example:
High-Throughput Computing System

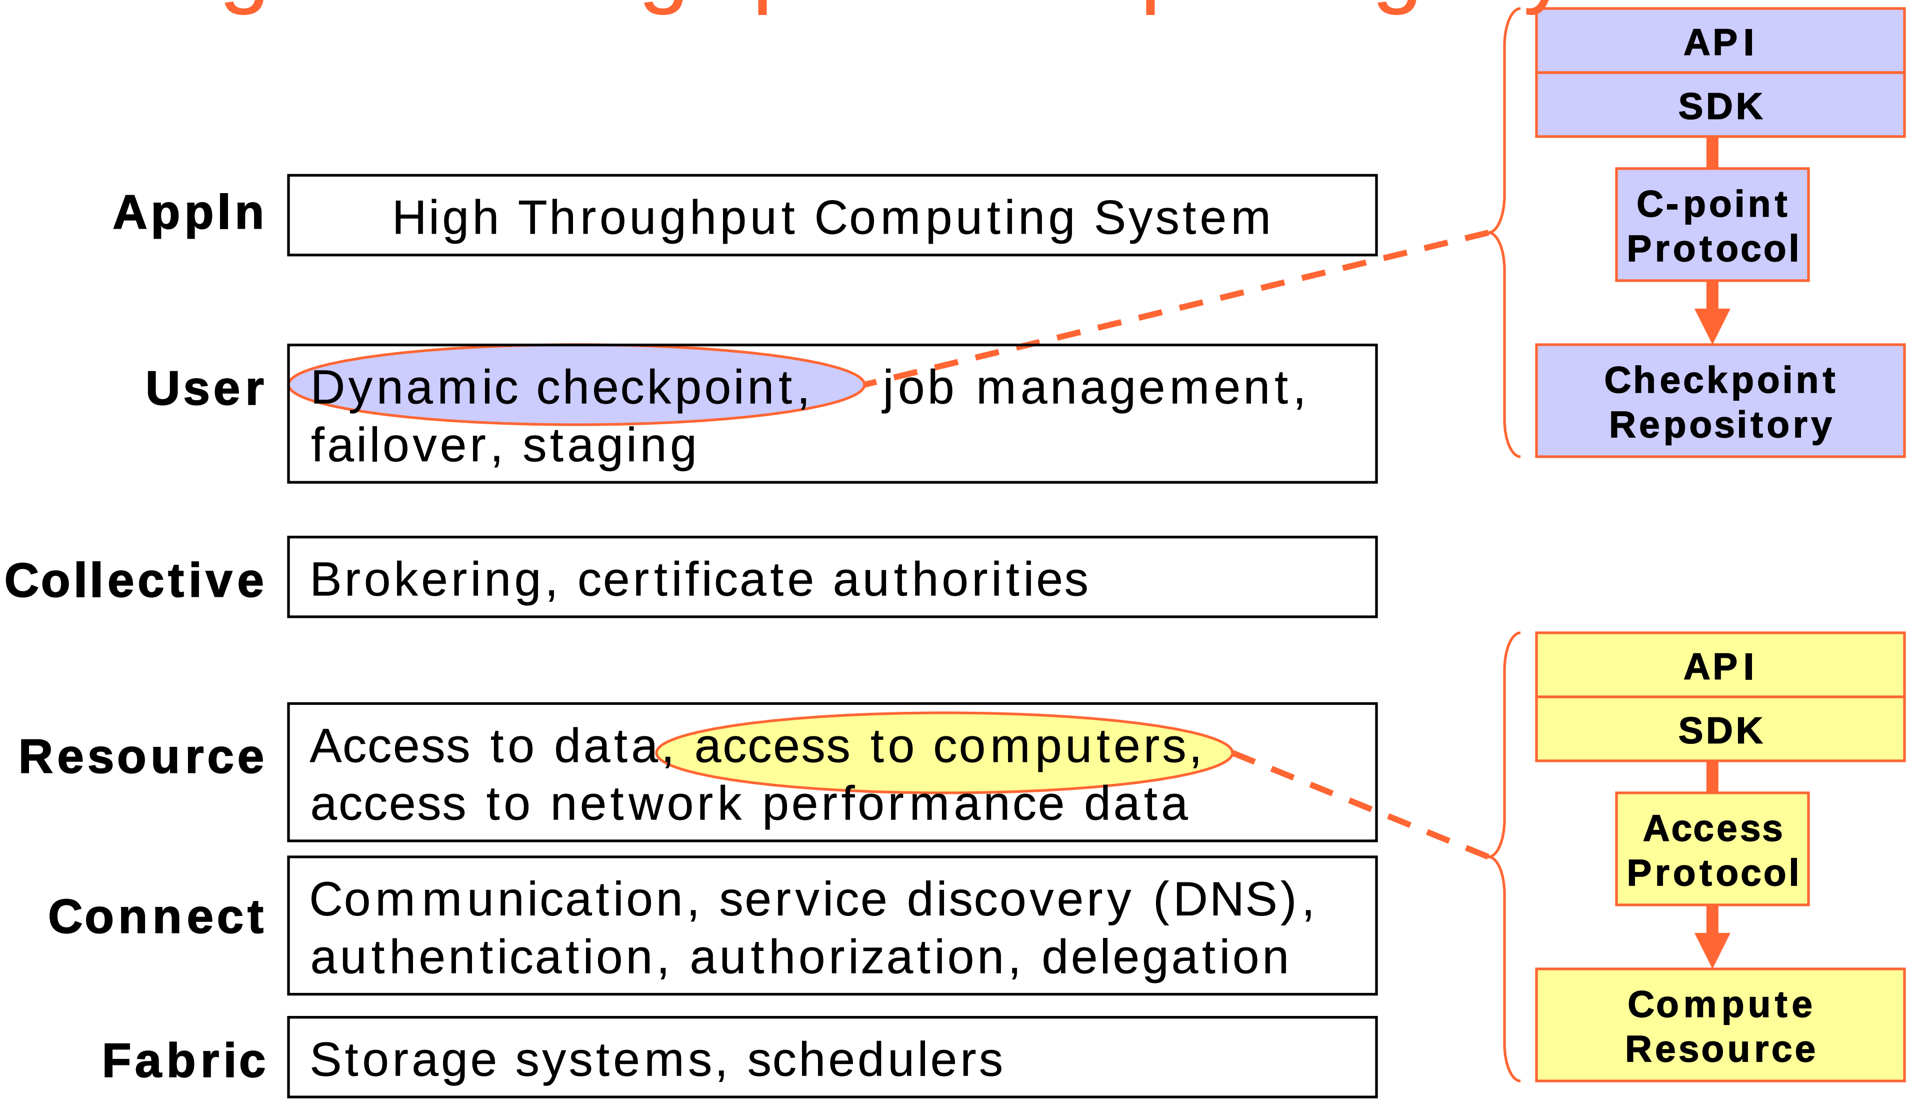

Data Grid R&D

- "Enable a geographically distributed community to pool their resources to perform sophisticated, computationally intensive analyses on Petabytes of data"
- Specific issues addressed
  - Network QoS for bulk data transfer
  - Replica management technologies
  - High-speed data transfer protocols & tools
  - Community-based access control
- Applications in climate and physics

Network Quality of Service Research: Overall Picture

- Secure, policy-driven bandwidth allocation for high-end applications
- Immediate and advance reservations
- Co-reservation and co-allocation of multiple resources for end-to-end flows
- All supported in a modular architecture widely used for Grid apps (Globus Toolkit)
- Experience with differentiated services
- Future opportunities in all-optical context
- Collaborative effort with LBNL

The GARNET QoS Research Testbed

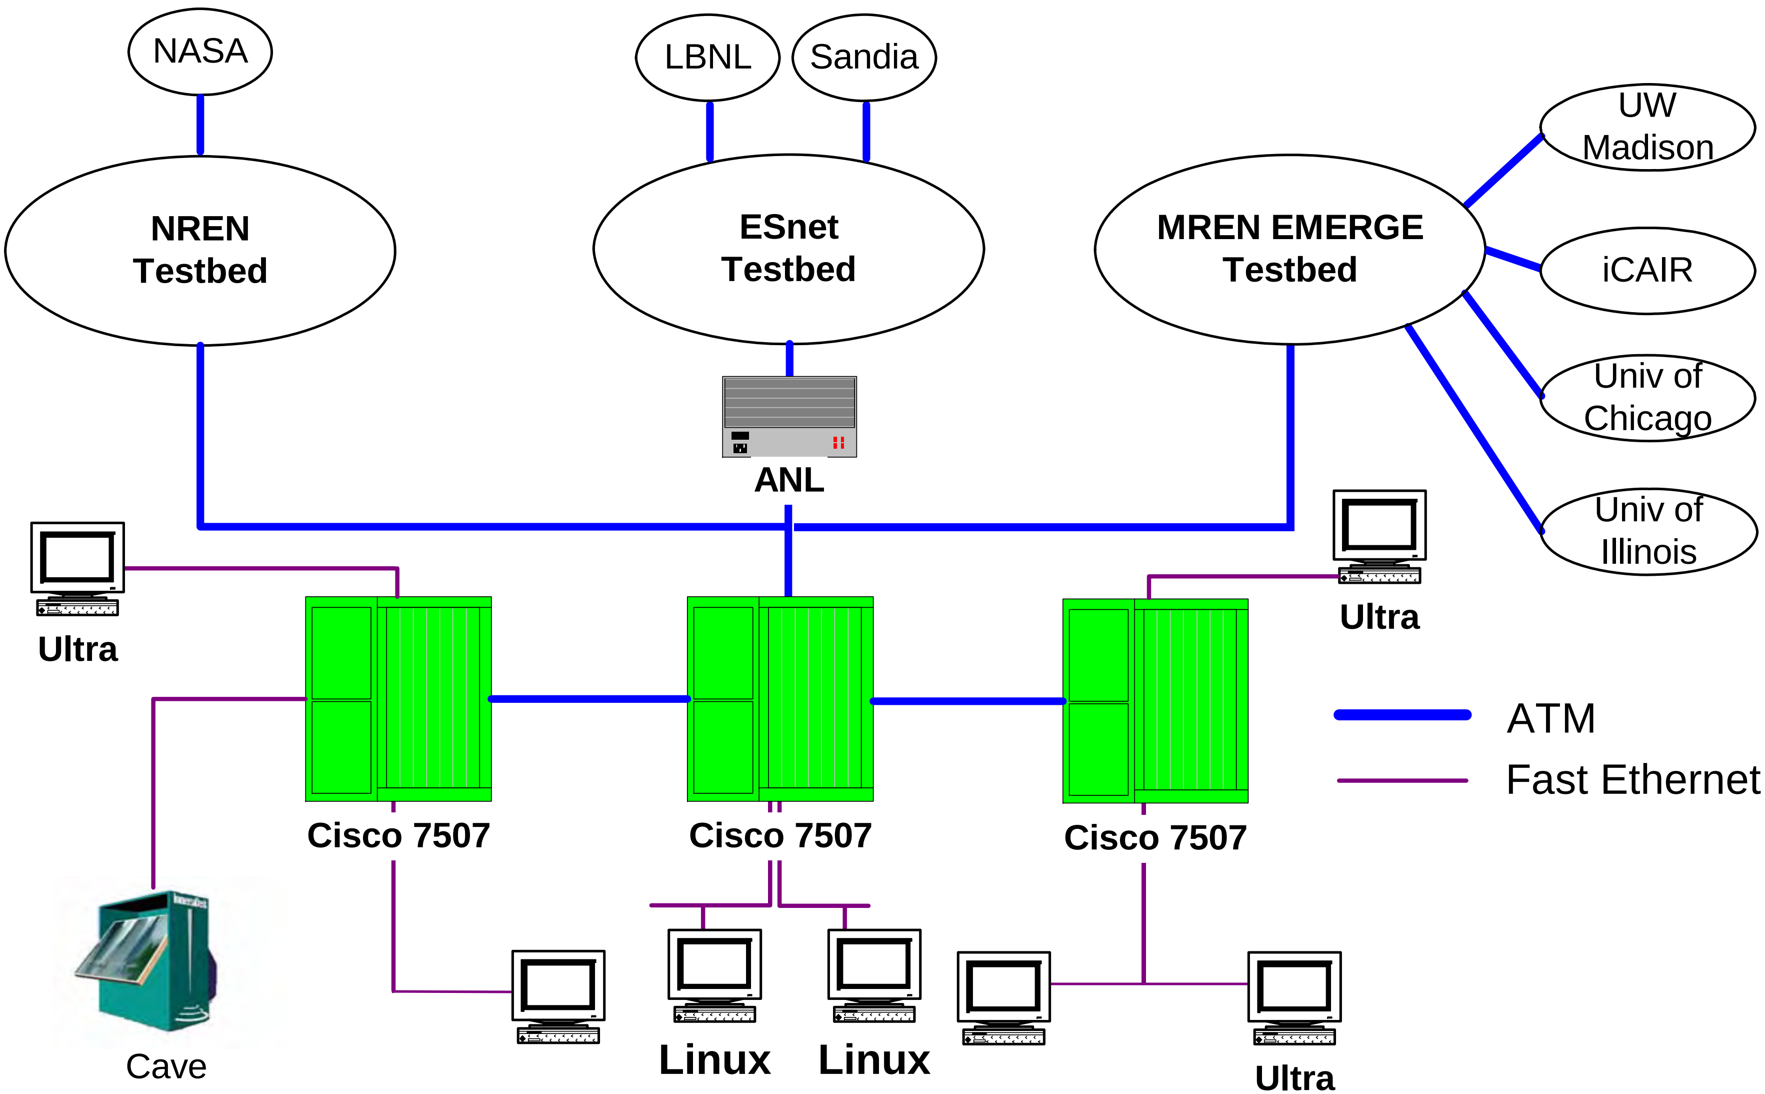

Bulk Transfer (LAN)

When a reservation ends, the bulk-transfer speeds up

When a reservation begins, the bulk-transfer backs off

The competitive UDP traffic never interferes

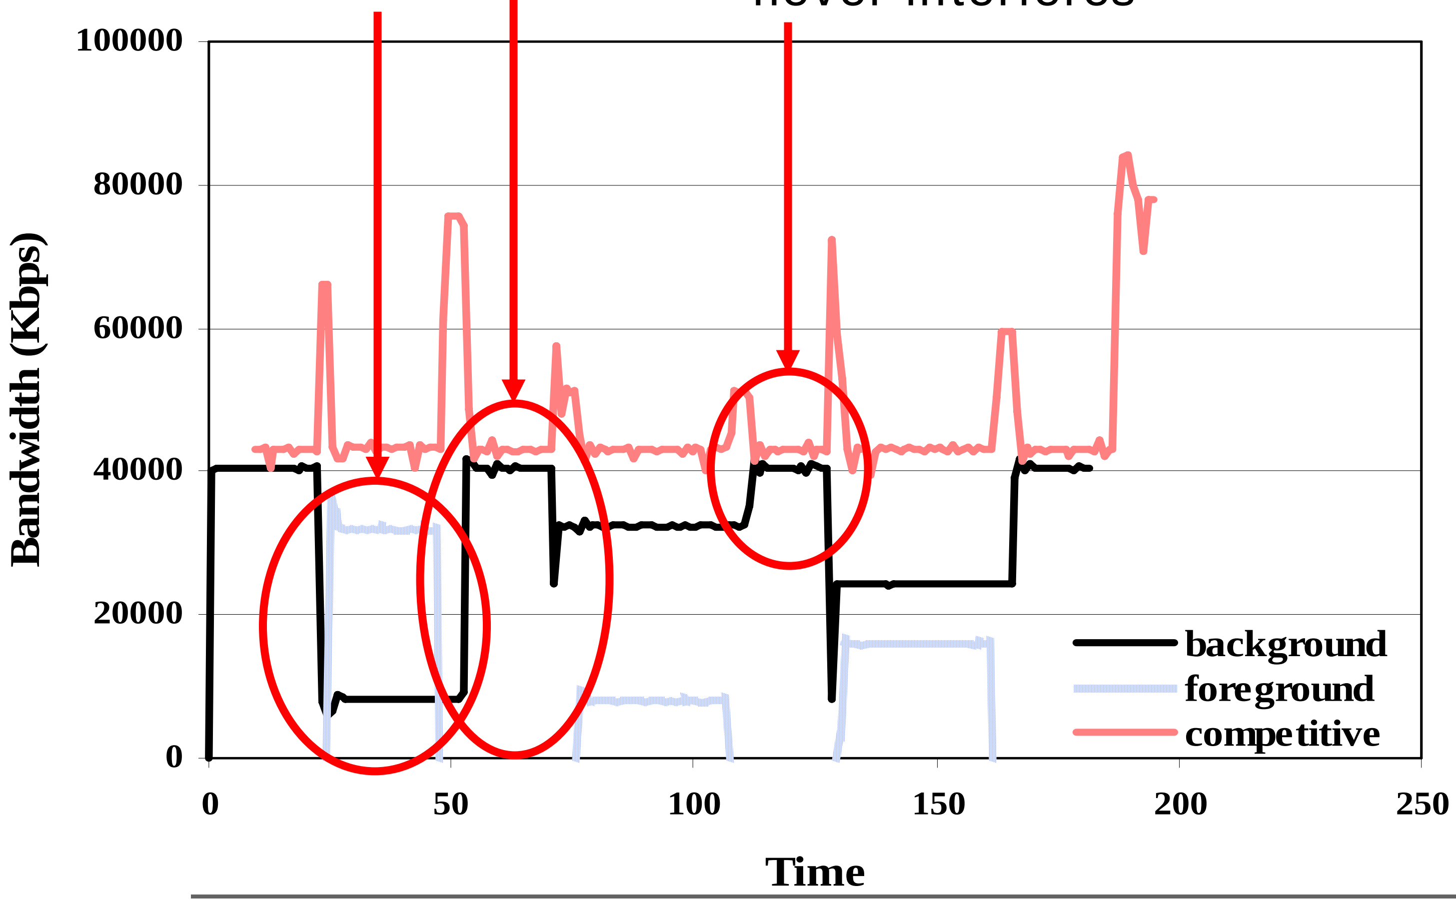

Bulk Transfer (WAN)

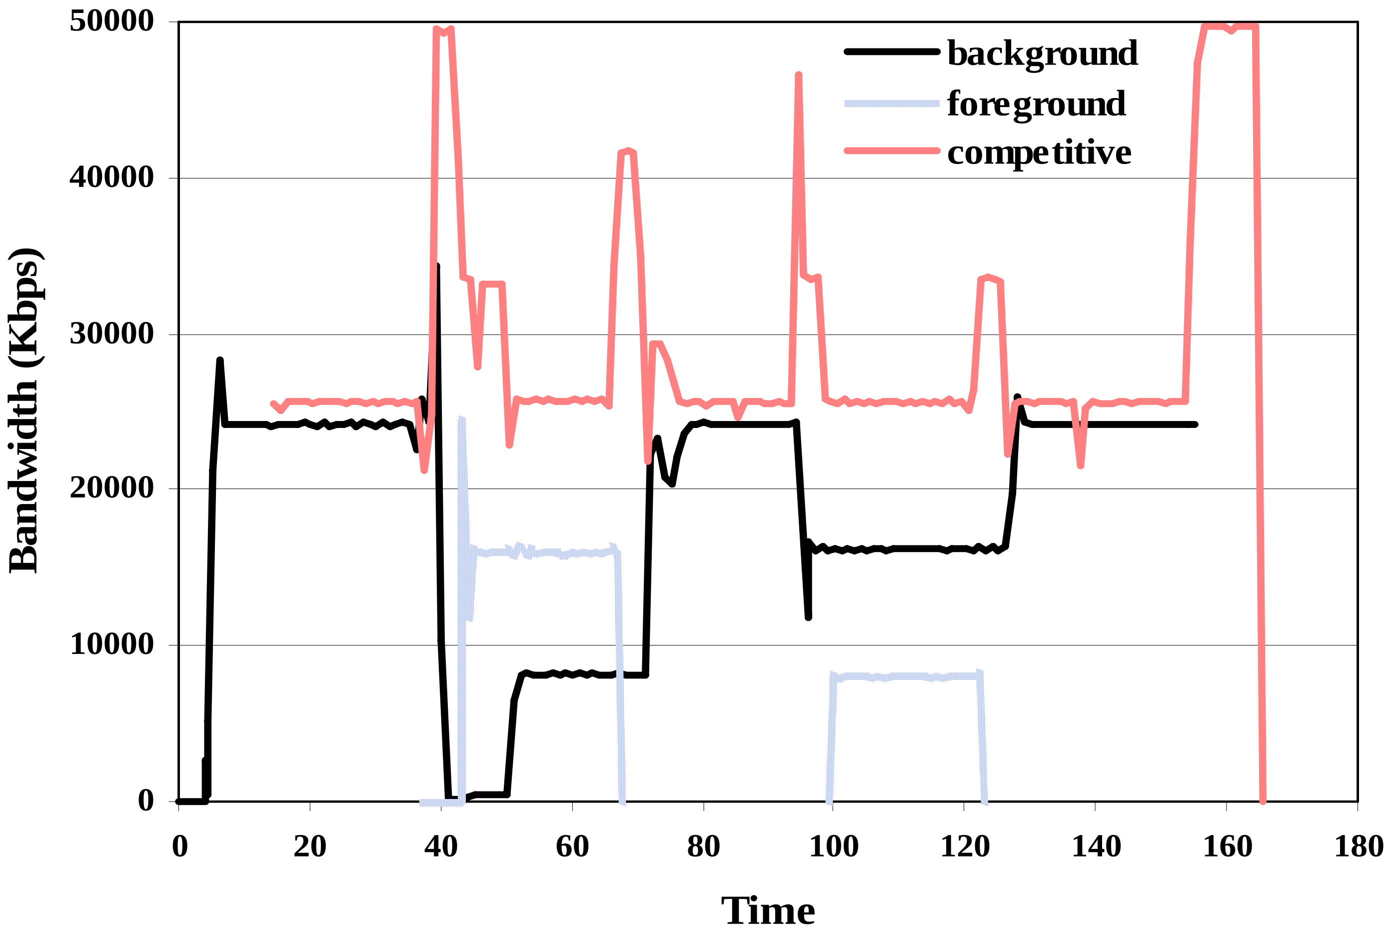

Data Grid Architecture

Appln	Discipline-Specific Data Grid Application
User	Coherency control, replica selection, task management, virtual data catalog, virtual data code catalog, …
Collective	Replica catalog, replica management, co-allocation, certificate authorities, metadata catalogs,
Resource	Access to data, access to computers, access to network performance data, …
Connect	Communication, service discovery (DNS), authentication, authorization, delegation
Fabric	Storage systems, clusters, networks, network caches, …

Globus Data-Intensive Services Architecture

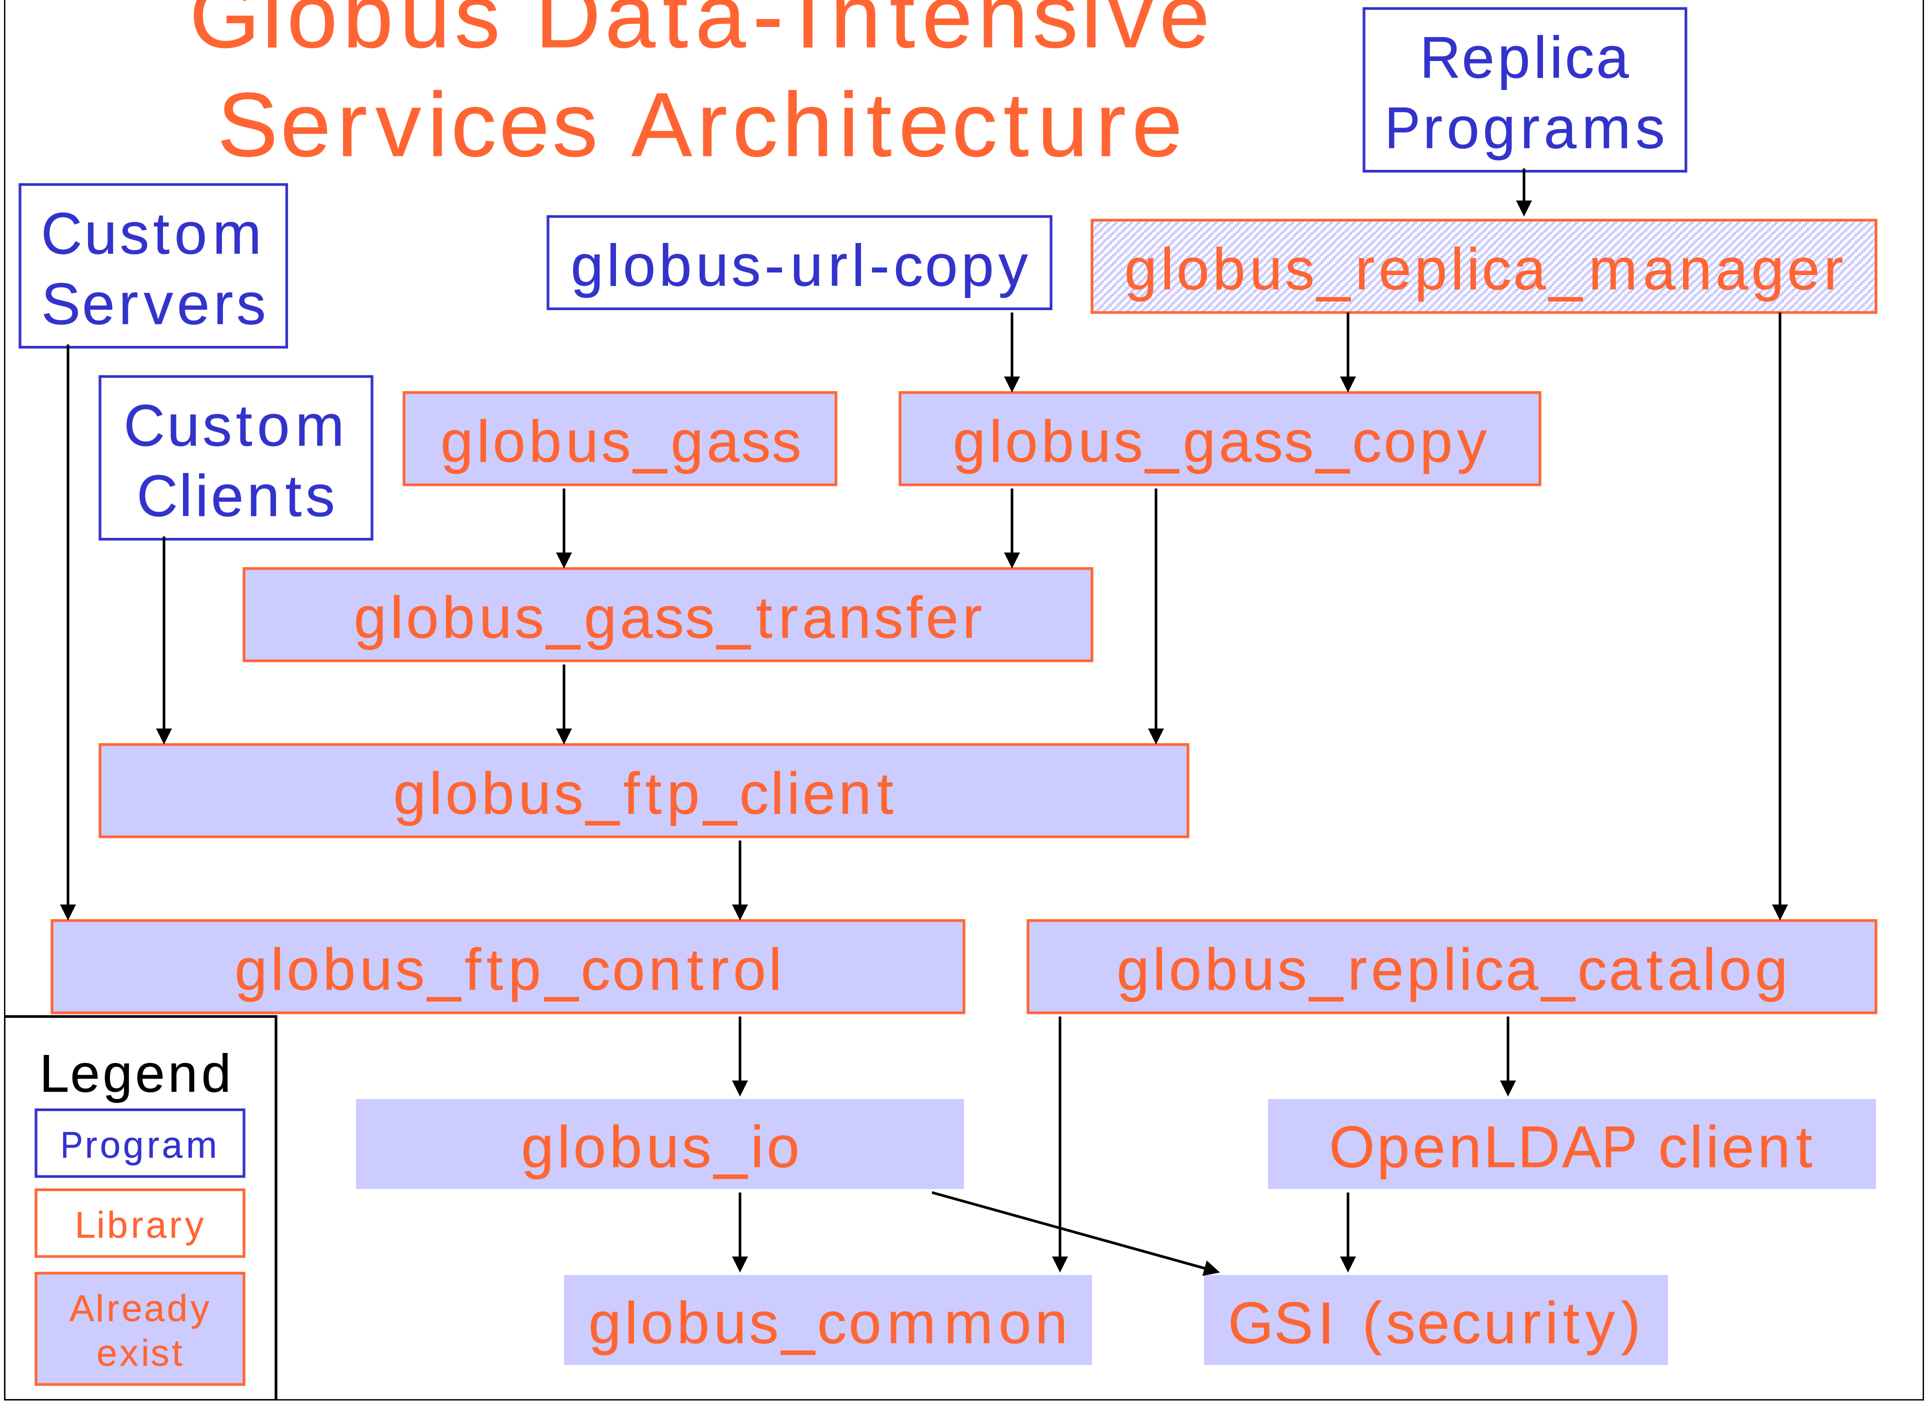

High-Level View of Earth System Grid: A Model Architecture for Data Grids

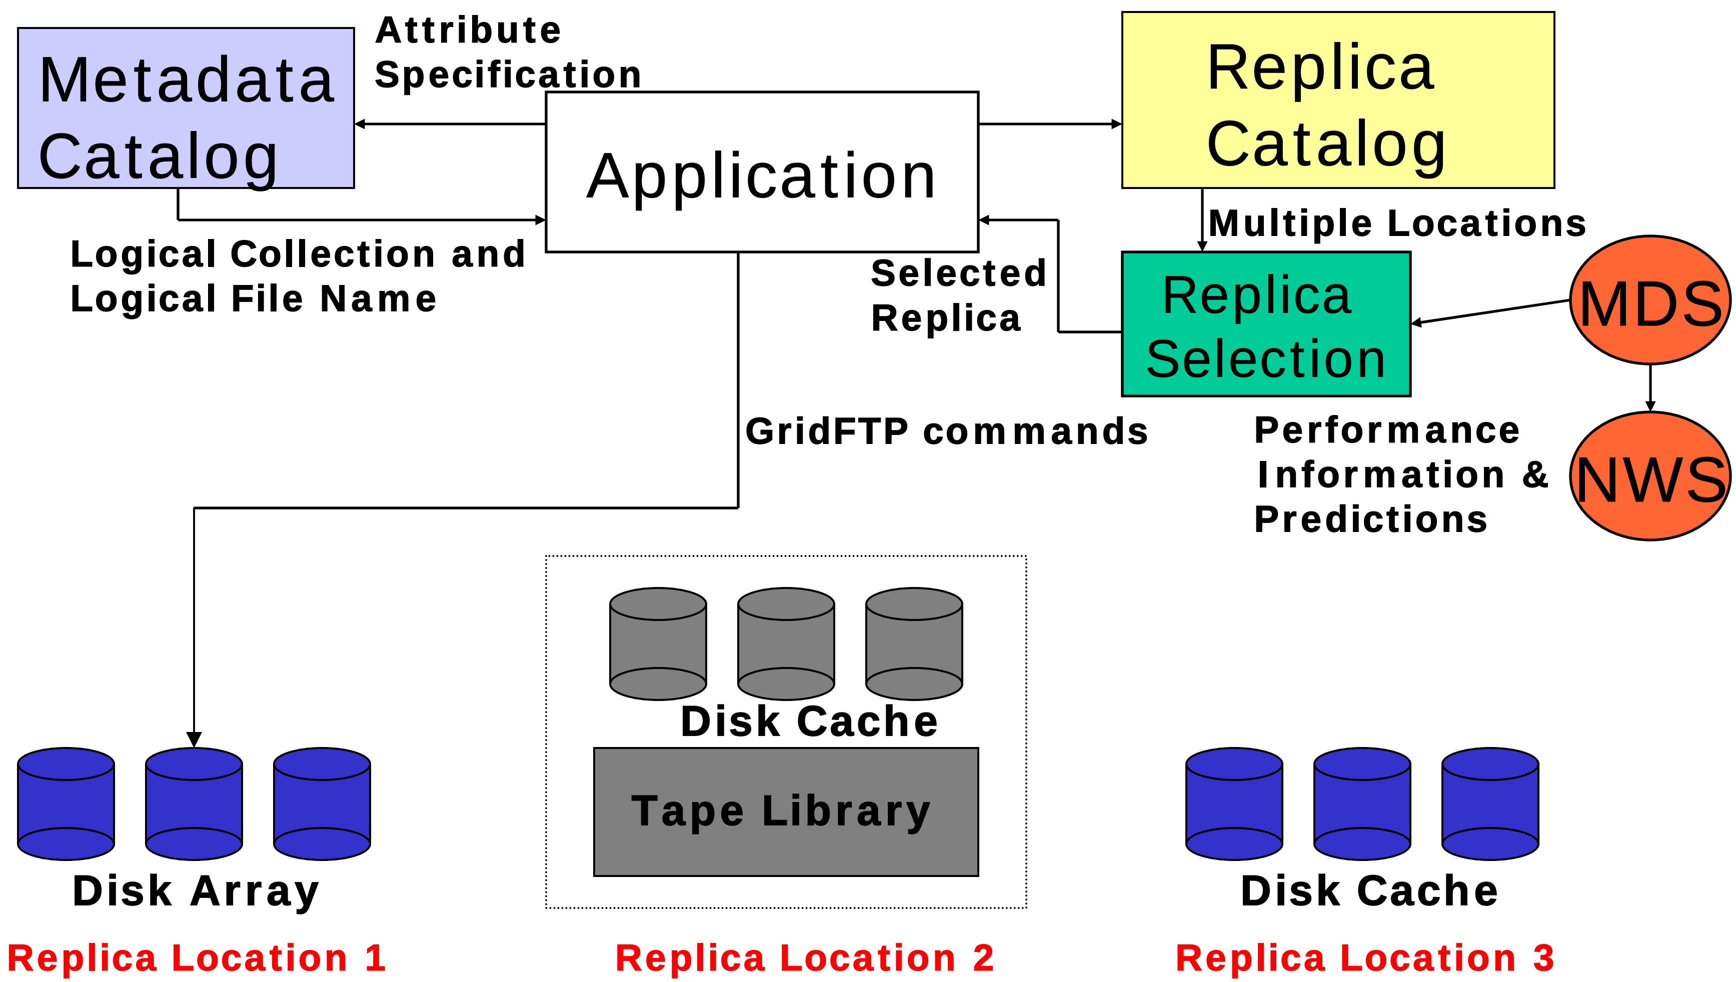

Future Directions

- "Grid computing" represents a significant success story for DOE research and opportunity for turbocharging DOE science
- Timely to take things to the next level, by
  - Aggressive programs focused on Grid-enabling key DOE applications (climate, physics, combustion, etc.)
  - New efforts focused on security, next-gen optical technologies, Grid tools, etc.
  - Infrastructure: faster networks, protocols, services (ESnet -> ESgrid?)

Abbreviated Acknowledgments

- Globus project Co-PI: Carl Kesselman
- Other Globus principals include: Steve Tuecke, Ann Chervenak, Bill Allcock, Gregor von Laszewski, Lee Liming, Steve Fitzgerald
- Access Grid: Rick Stevens & many others
- DOE Grid collaborators: Arie Shoshani, Reagan Moore, Miron Livny, Bill Johnston, Brian Tierney, others
- ESG collaborators: Dean Williams & others
- PPDG PIs: Richard Mount, Harvey Newman

Future Directions

- "Grid computing" represents a significant success story for DOE research and opportunity for turbocharging DOE science
- Timely to take things to the next level, by
  - Aggressive programs focused on Grid-enabling key DOE applications (climate, physics, combustion, etc.)
  - New efforts focused on security, next-gen optical technologies, Grid tools, etc.
  - Infrastructure: faster networks, protocols, services (ESnet -> ESgrid?)

For More Information

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