The Data Grid:

An Architecture for Distributed Management of Large Scientific Data Sets

Ann Chervenak
Carl Kesselman

Ian Foster
Chuck Salisbury
Steve Tuecke

Information Sciences Institute,
University of Southern California

Argonne National Laboratory

Overview

- Target Environment
- Design Principles
- Grid Services
  - Storage systems
  - Metadata
  - Management of replicated files
- Implementation

Data Grid Environment

- Scientific applications
  - Global climate change, High energy physics
- Computationally demanding
- Large data sets and archives
  - Terabytes, eventually petabytes
  - Raw and derived data
- Geographically dispersed users and resources
  - Data replication for enhanced performance
- Broad range of capabilities and resources
  - Networks, systems, storage, and applications

Building a Data Grid: Building Blocks

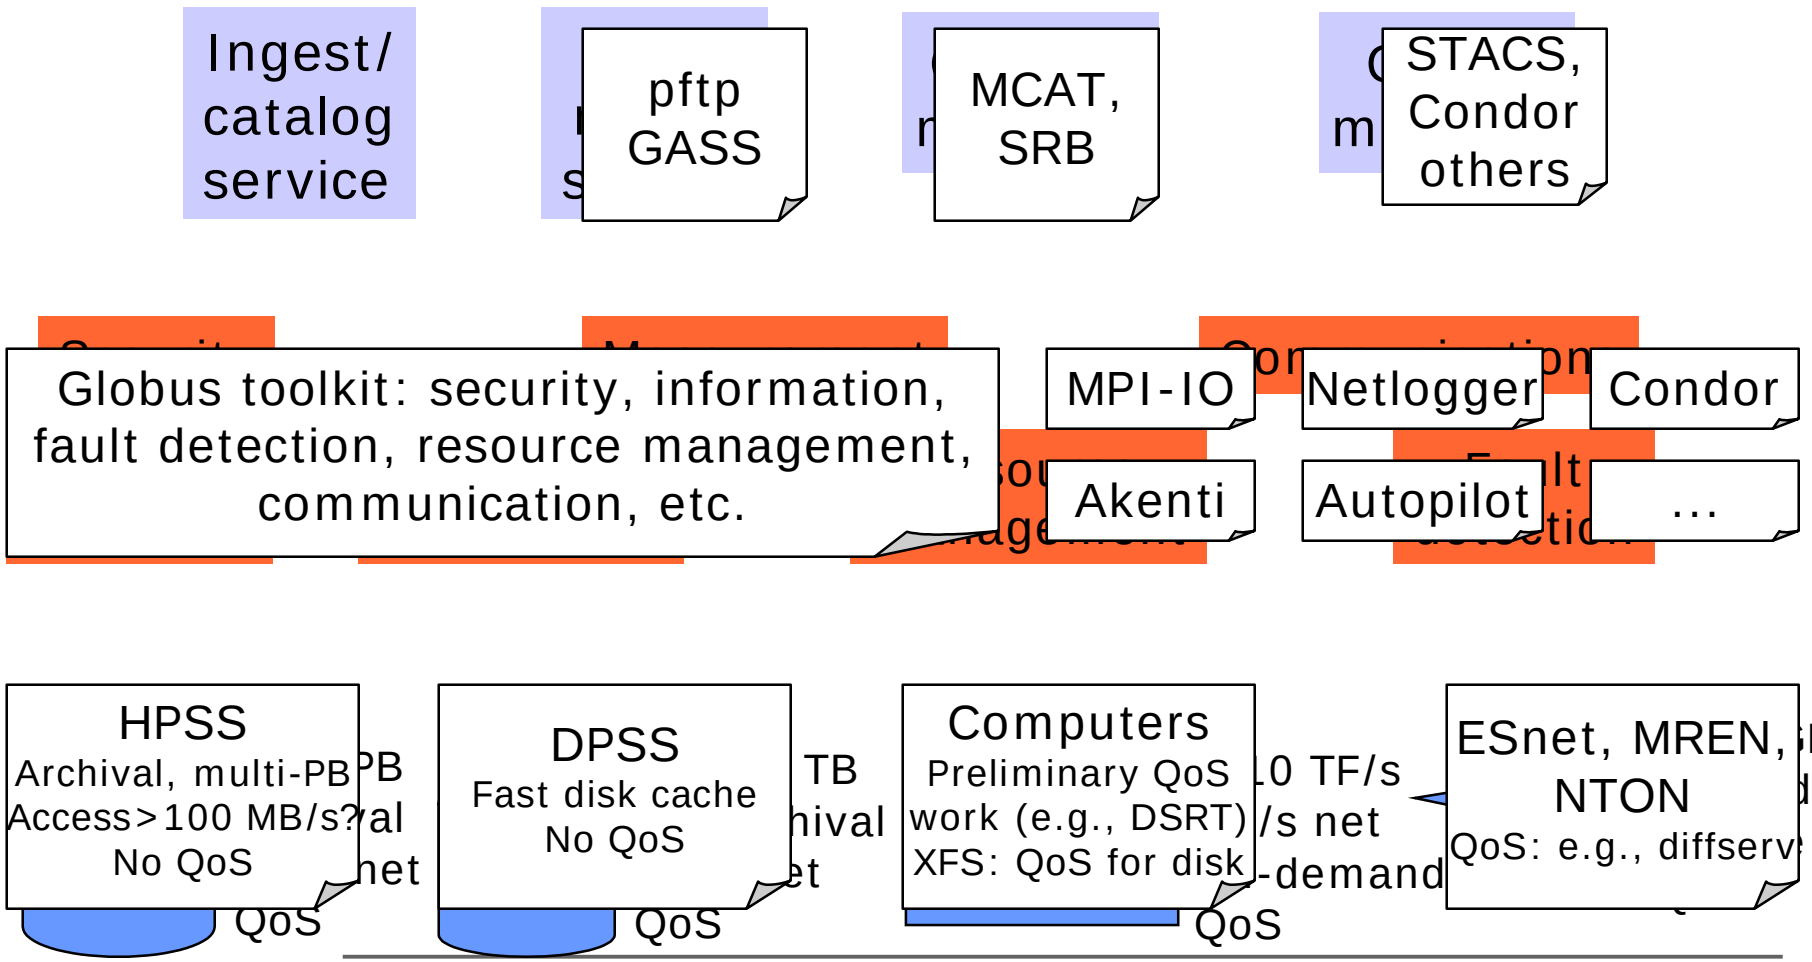

Data Grid Objectives

- Integrate heterogeneous data archives into a distributed data management “grid”
- Identify services for high performance, distributed, data intensive computing

Design Principles

- Mechanism Neutrality
  - Support heterogeneous systems
- Policy Neutrality
  - User / local decision making and control
- Compatibility with Computational Grid
  - Integration of storage and computation
- Uniformity of Information Infrastructure
  - Data model and interface for metadata

Data Grid Services

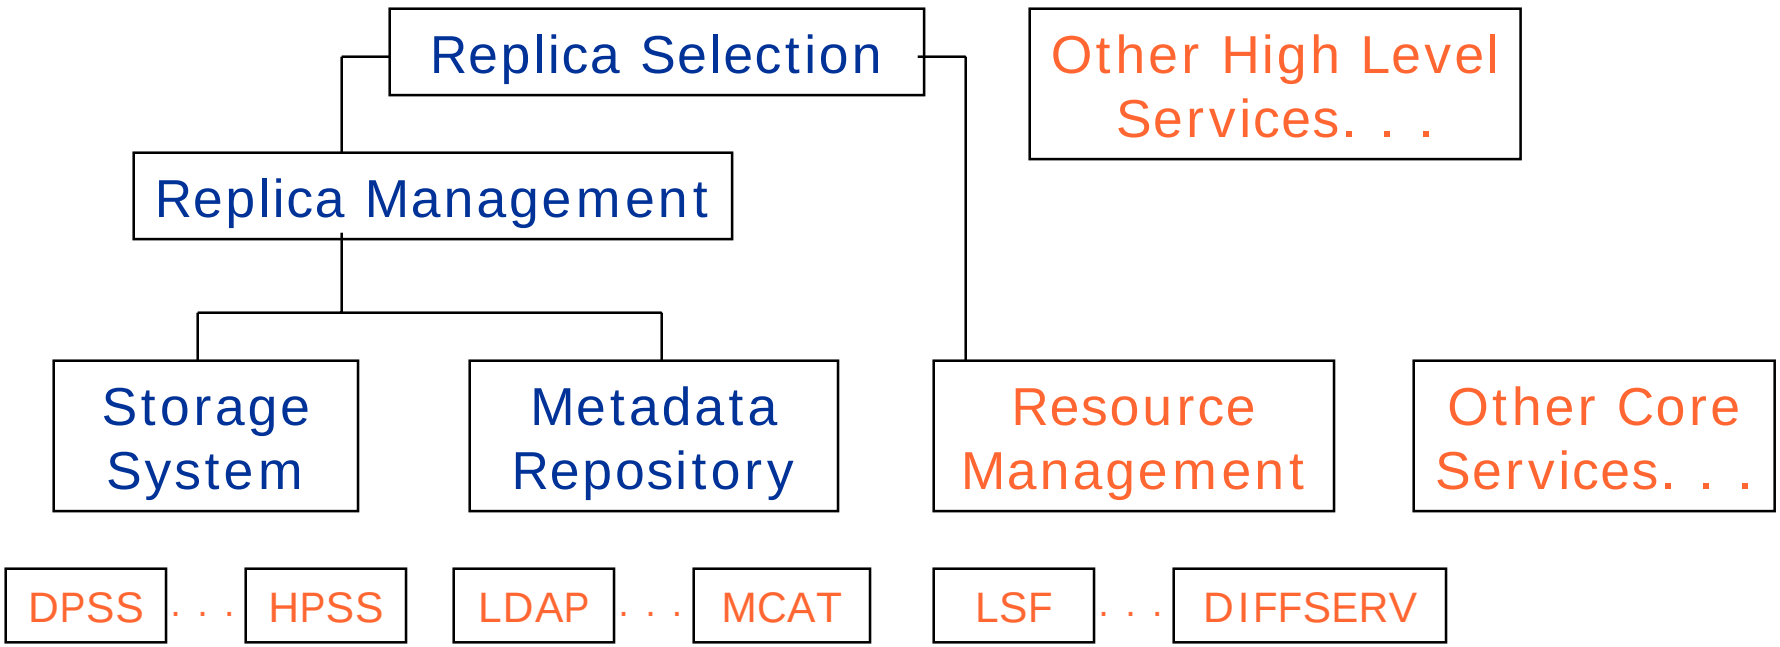

Data Access Service

- Uniform access to heterogeneous systems
  - remote: *e.g.* DPSS, HTTP, FTP, HPSS
  - local: *e.g.* UNIX
- High performance data movement over WANs
  - Third party transfer
- Data extraction and filtering functions
- Access to data is subject to global and local policy constraints

Metadata Access Service

- Uniform treatment for all metadata
  - Grid components
  - Application-related metadata
  - Storage system characteristics
  - Relationships between data items
- Uniform access to metadata
  - LDAP protocol
- Uniform storage structure
  - LDAP hierarchical structure for distribution, replication, referral services

Replica Management

- Collections contain related files
- Logical files describe replicated physical files
- Services for managing replicated file instances
  - Create / delete
  - Schedule / manage data transfer
  - Register in the replica catalog
  - Metadata display

Replica Selection

- User can optimize access characteristics
  - Grid structure and performance
  - Storage system and file characteristics
- Intelligent scheduling to determine appropriate replica, site for (re)computation, etc.

Climate Data Scenario

"How do midwest flood frequencies under 2xCO_2 scenario compare with historical data?"

Query manager

"Access datasets A, B; run A->meso->hydro; compare result with B"

File access service

Historical data archive

Resource manager

Cache

Simulation data archive

DPSS

HPSS

Analysis engine

meso → hydro → compare

Current Activity

- Ongoing collaborations
  - Climate
  - High Energy Physics
- Storage API for uniform access to data
  - API specification document
  - Prototype code for HTTP, FTP, DPSS
- Replica management
  - Replica catalog based on LDAP
  - API and GUI tools for catalog access
- Quality of Service implementation

the globus project
www.globus.org

Replica Management

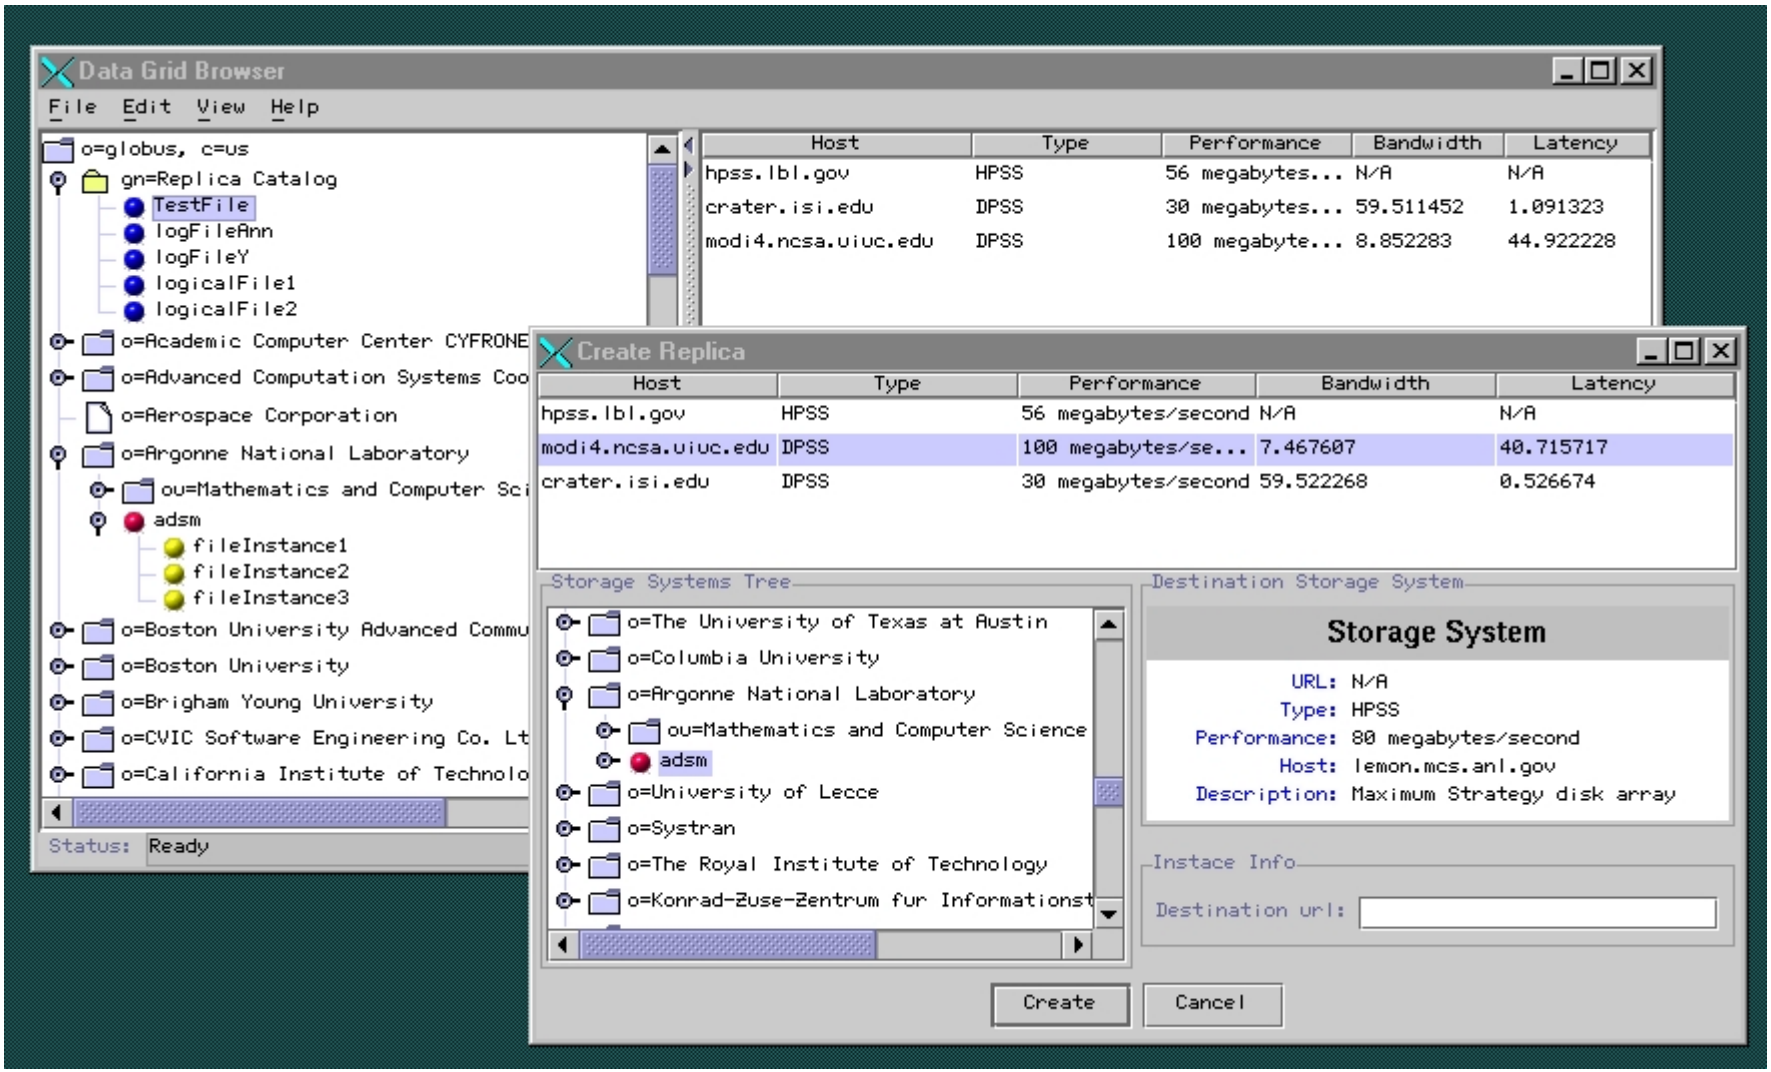

Quality of Service

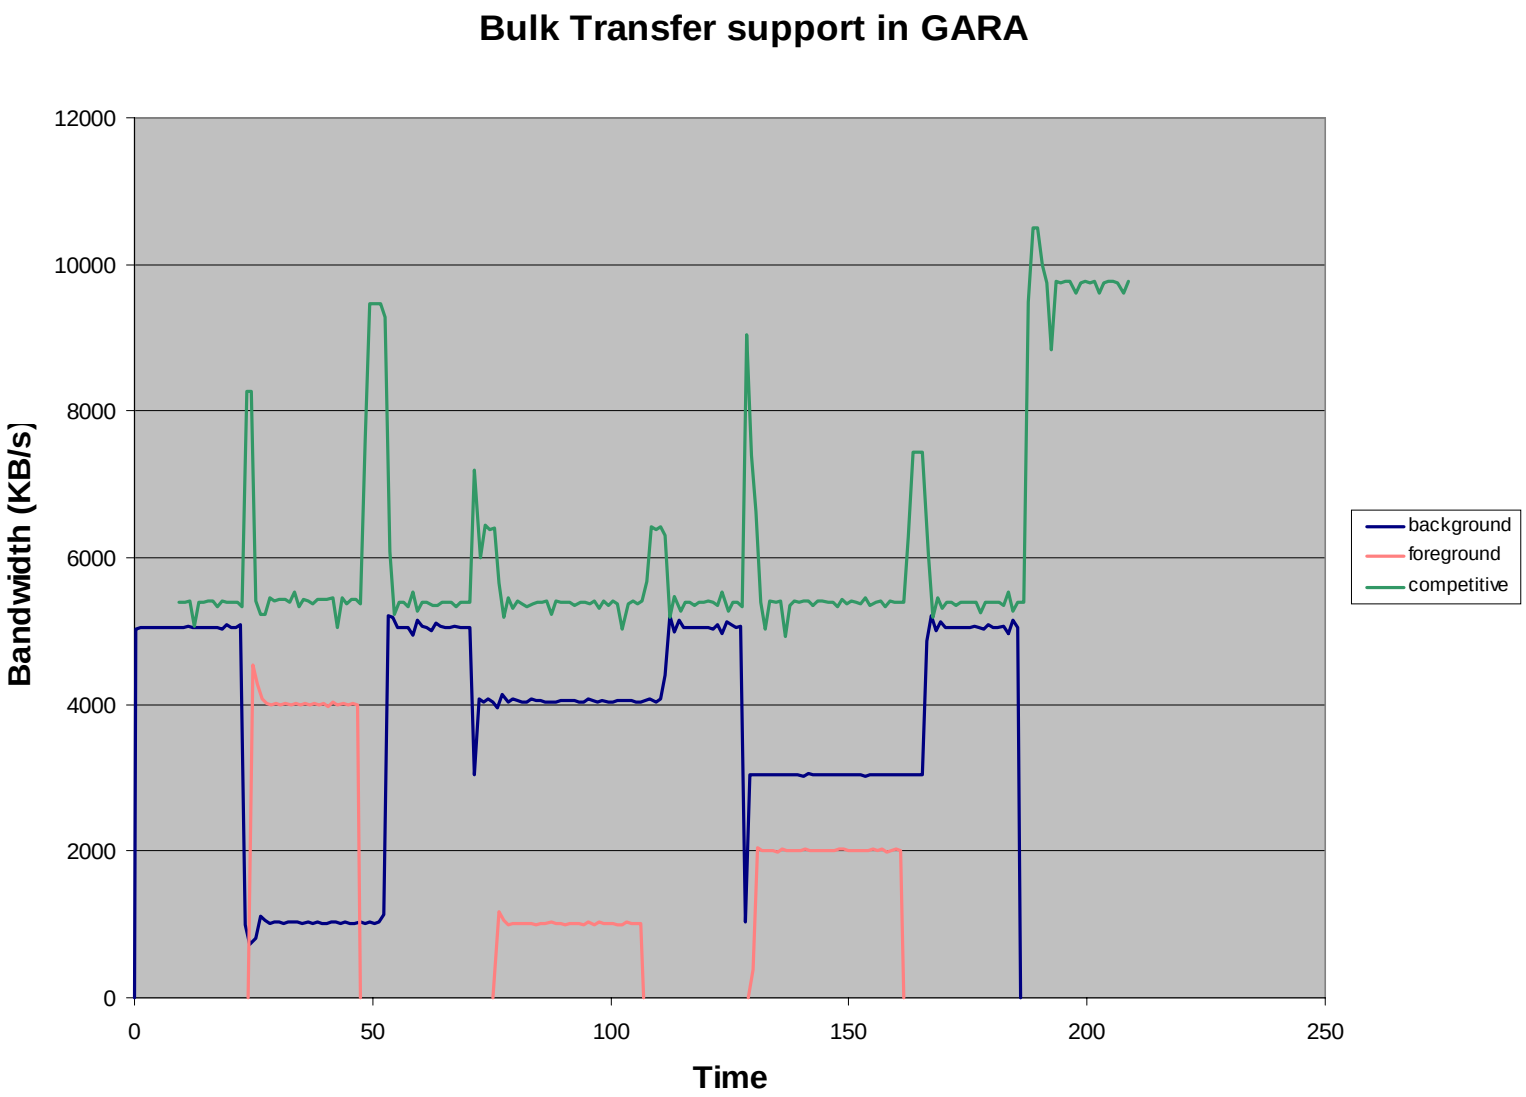

Planned Activity

- Data Access
  - Integrated quality of service, security
  - Performance enhancements for networking
- Performance guarantees for the Data Grid
- Automatic operation of the Data Grid
  - Agent technologies used for distributed data replication, selection, and analysis
- Integrated CPU scheduling
  - Server-side data reduction, affinity scheduling


DPSS
Visualization
LBNL
SCinet '99
ASCI
ANL
DPSS
DPSS
DPSS
ADSM
Visualization
DPSS
DPSS
SRB
Unitree