Building an Open Grid

Ian Foster

Argonne National Laboratory
University of Chicago

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

2nd Access Grid Retreat, Argonne, April 15, 2003

the globus project™
www.globus.org

Overview

- Grid: why and what
- Evolution of Grid technology
  - Open Grid Service Architecture
- Future directions
  - Towards lightweight VOs: dynamic trust relationships
  - Towards global knowledge communities: virtual data and dynamic workspaces

the globus project™
www.globus.org

Why the Grid?
(1) Revolution in Science

- Pre-Internet
  - Theorize &/or experiment, alone or in small teams; publish paper

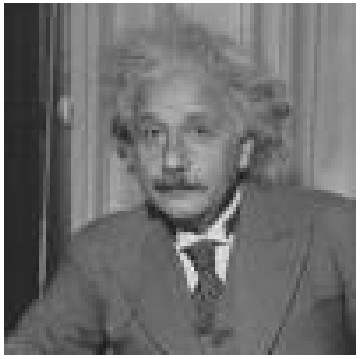

- Post-Internet
  - Construct and mine large databases of observational or simulation data
  - Develop simulations & analyses
  - Access specialized devices remotely
  - Exchange information within distributed multidisciplinary teams

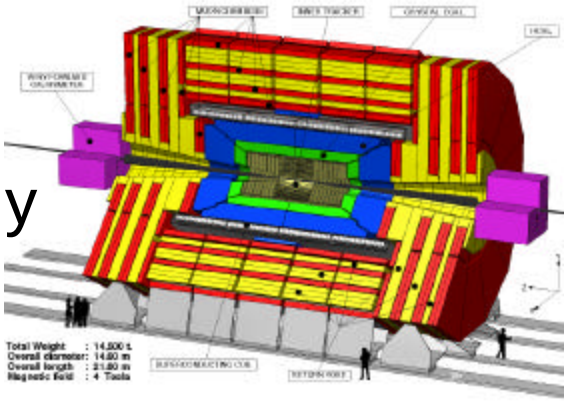

the globus project™
www.globus.org

Why the Grid?
(2) Revolution in Business

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

the globus project™
www.globus.org

New Opportunities Demand New Technology

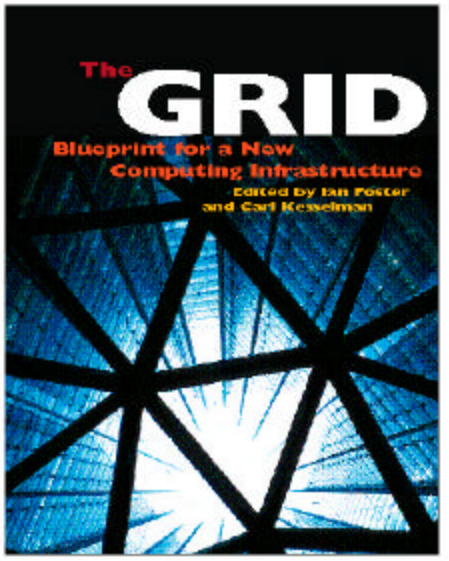

"Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations"

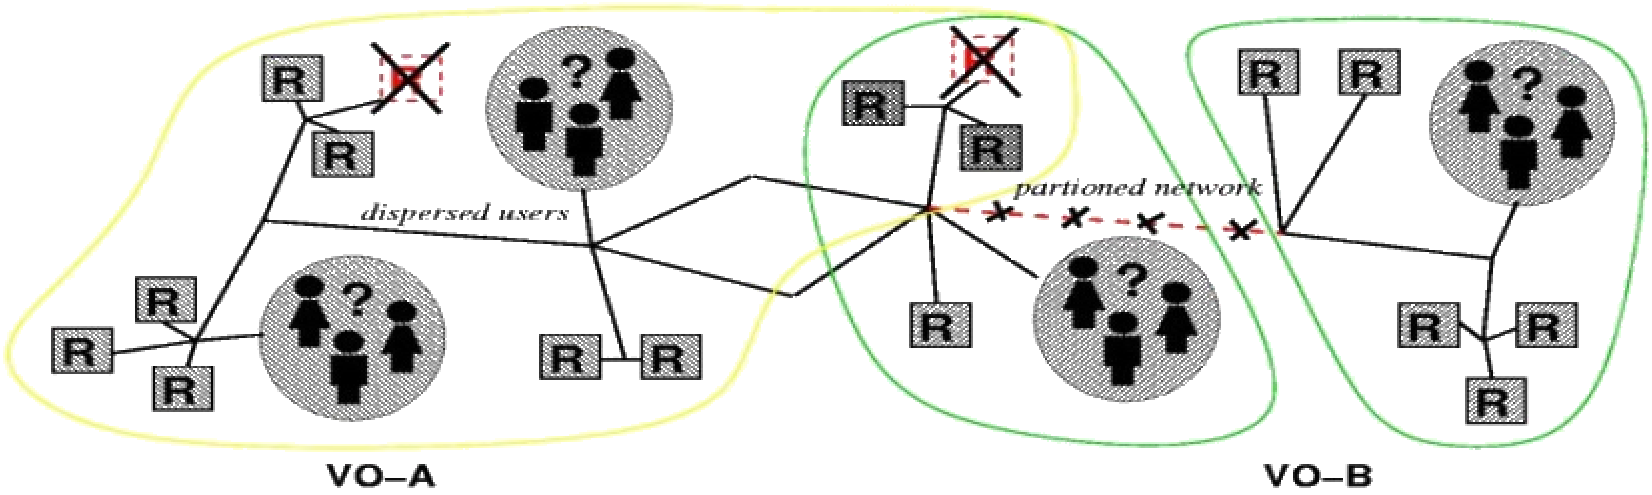

"*When the network is as fast as the computer's internal links, the machine disintegrates across the net into a set of special purpose appliances*" (George Gilder)

foster@mcs.anl.gov

the globus project™
www.globus.org

Grid Communities & Technologies

- Yesterday
  - Small, static communities, primarily in science
  - Focus on sharing of computing resources
  - Globus Toolkit as technology base
- Today
  - Larger communities in science; early industry
  - Focused on sharing of data and computing
  - Open Grid Services Architecture emerging
- Tomorrow
  - Large, dynamic, diverse communities that share a wide variety of services, resources, data
  - New issues: Trust, distributed RM, knowledge

the globus project™
www.globus.org

NSF TeraGrid

- NCSA, SDSC, Argonne, Caltech
- Unprecedented capability
  - 13.6 trillion flop/s
  - 600 terabytes of data
  - 40 gigabits per second
  - Accessible to thousands of scientists working on advanced research
- www.teragrid.org

foster@mcs.anl.gov

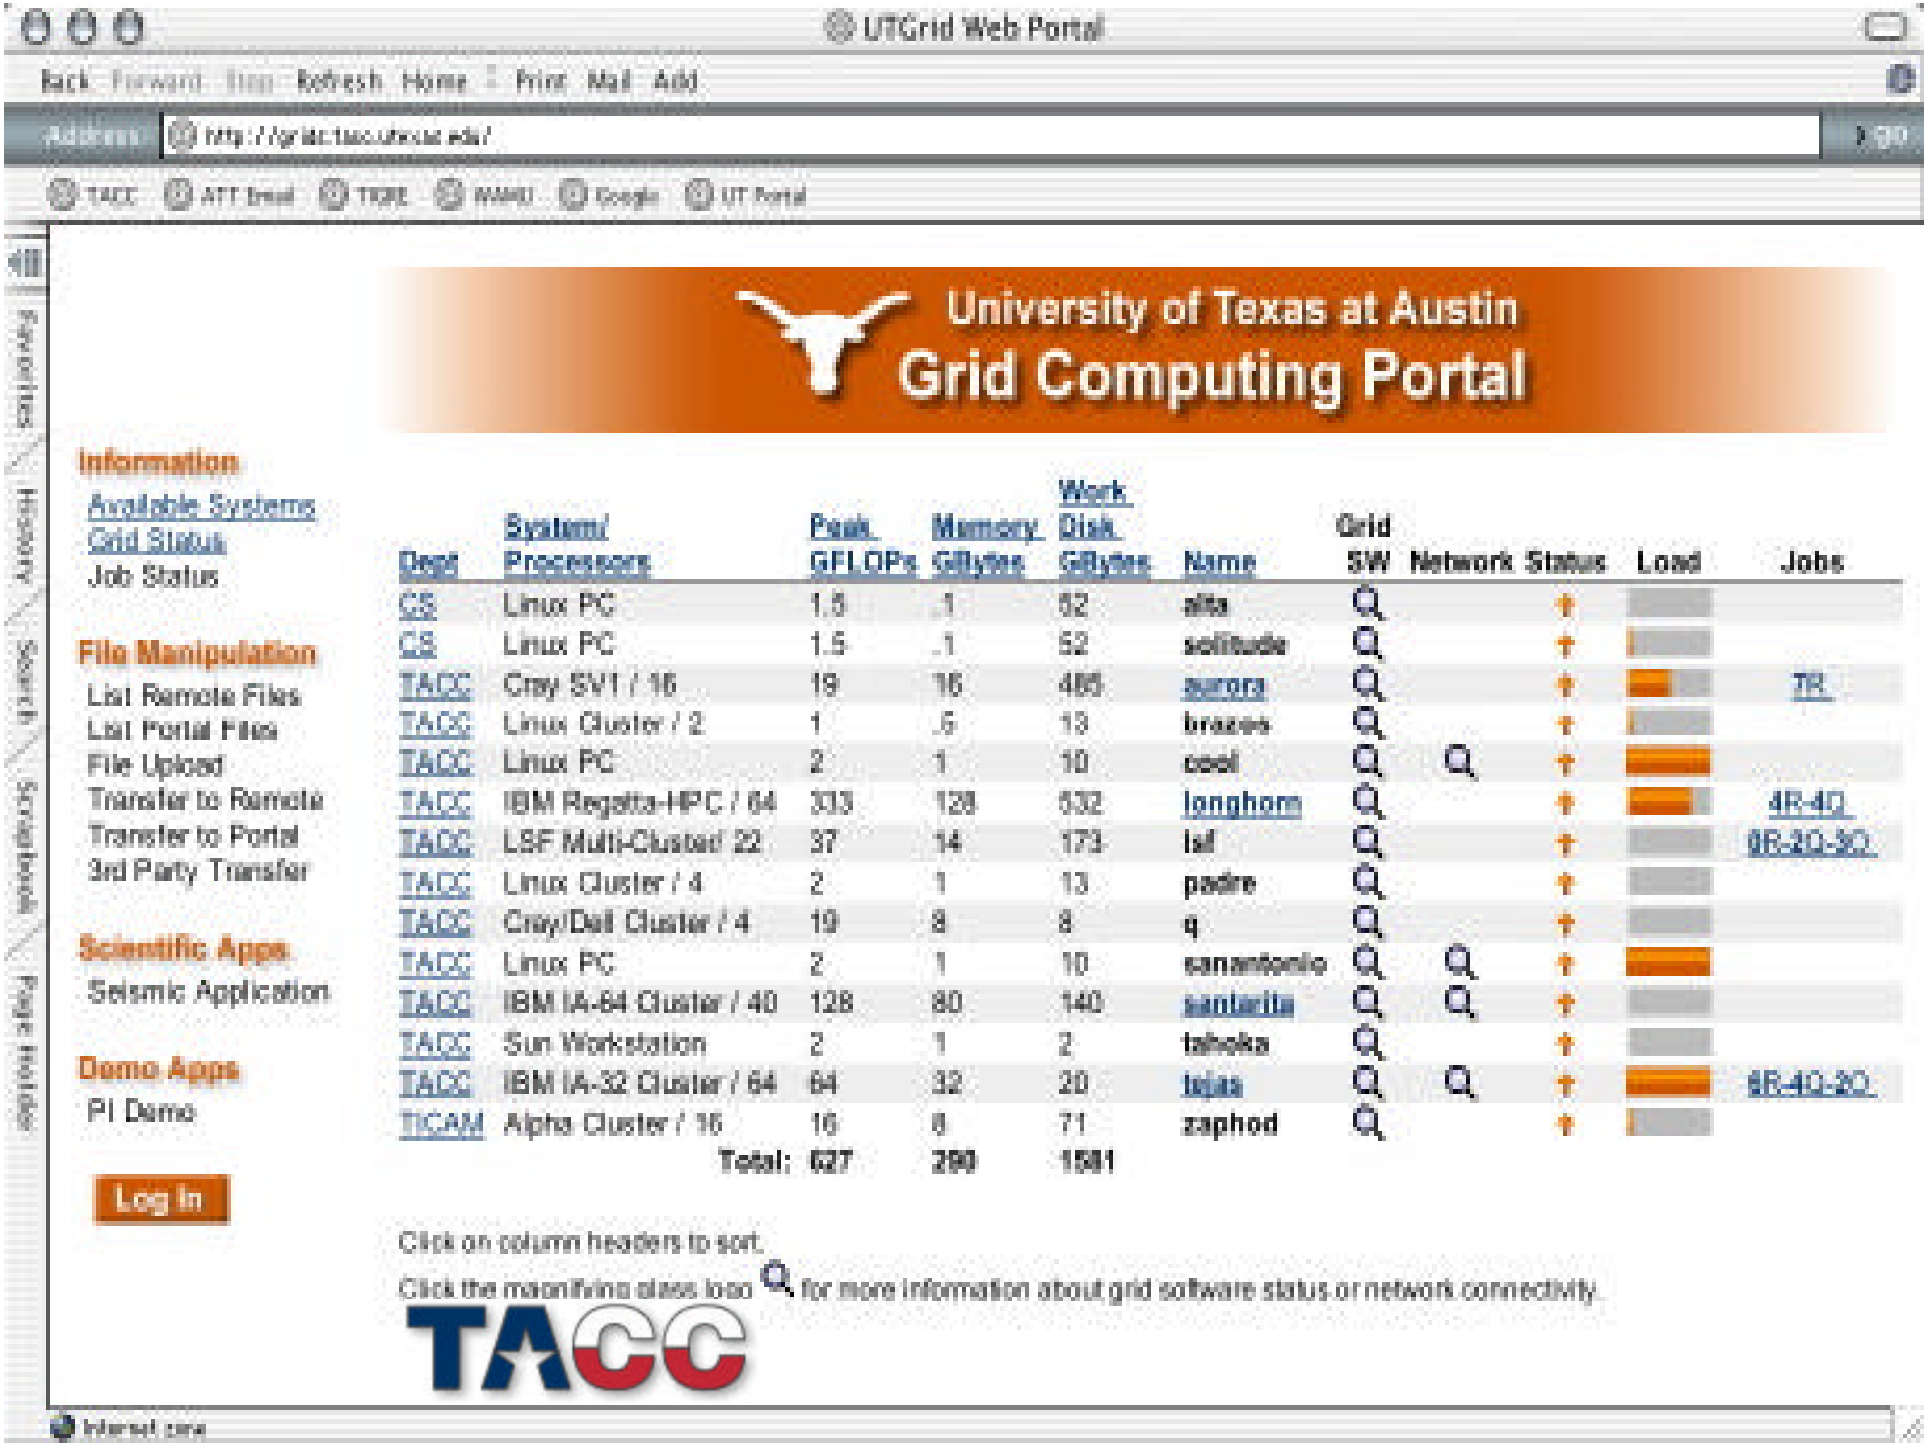
UTGrid Web Portal
Back Forward Stop Refresh Home Print Mail Add
Address http://grid.tacc.utexas.edu/
TACC ATT Email TIGRE WAMU Google UT Portal
Favorites History Search Scrapbook File Info
University of Texas at Austin
Grid Computing Portal
Information
Available Systems
Grid Status
Job Status
File Manipulation
List Remote Files
List Portal Files
File Upload
Transfer to Remote
Transfer to Portal
3rd Party Transfer
Scientific Apps
Seismic Application
Demo Apps
PI Demo
Log In
Dept | System/ Processors | Peak GFLOPs | Memory GBytes | Work Disk GBytes | Name | Grid SW | Network | Status | Load | Jobs
CS | Linux PC | 1.5 | .1 | 52 | alta
CS | Linux PC | 1.5 | .1 | 52 | solitude
TACC | Cray SV1 / 16 | 19 | 16 | 485 | aurora | TR
TACC | Linux Cluster / 2 | 1 | .5 | 13 | brazos
TACC | Linux PC | 2 | 1 | 10 | cool
TACC | IBM Regatta-HPC / 64 | 333 | 128 | 532 | longhorn | 4R-4Q
TACC | LSF Multi-Cluster/ 22 | 37 | 14 | 173 | lsf | 8R-2Q-3Q
TACC | Linux Cluster / 4 | 2 | 1 | 13 | padre
TACC | Cray/Dell Cluster / 4 | 19 | 8 | 8 | q
TACC | Linux PC | 2 | 1 | 10 | sanantonio
TACC | IBM IA-64 Cluster / 40 | 128 | 80 | 140 | santarita
TACC | Sun Workstation | 2 | 1 | 2 | tahoka
TACC | IBM IA-32 Cluster / 64 | 64 | 32 | 20 | tejas | 8R-4Q-2Q
TICAM | Alpha Cluster / 16 | 16 | 8 | 71 | zaphod
Total: 627 290 1581
Click on column headers to sort.
Click the magnifying glass icon for more information about grid software status or network connectivity.
TACC
Internet zone

Data Grids for High Energy Physics

- Enable international community of 1000s to access & analyze petabytes of data
- Harness computing & storage worldwide
- Virtual data concepts: manage programs, data, workflow
- Distributed system management

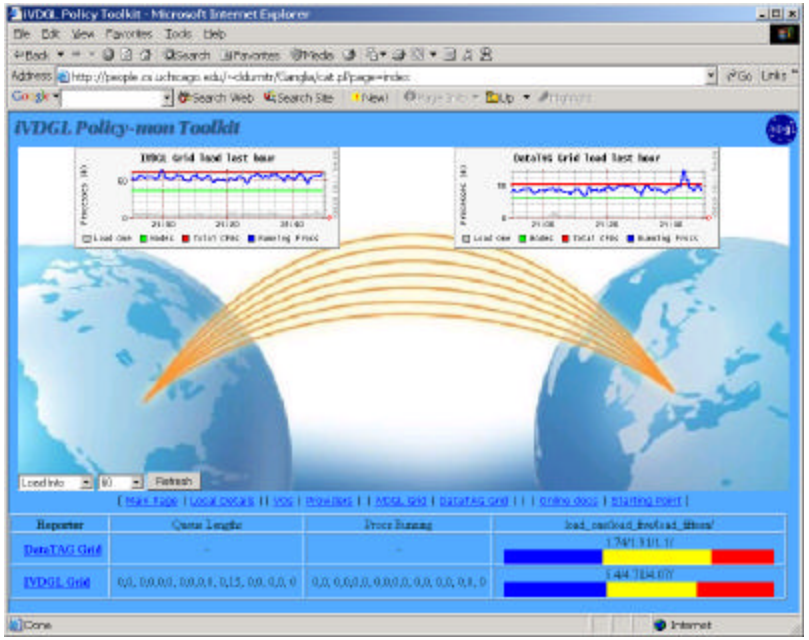

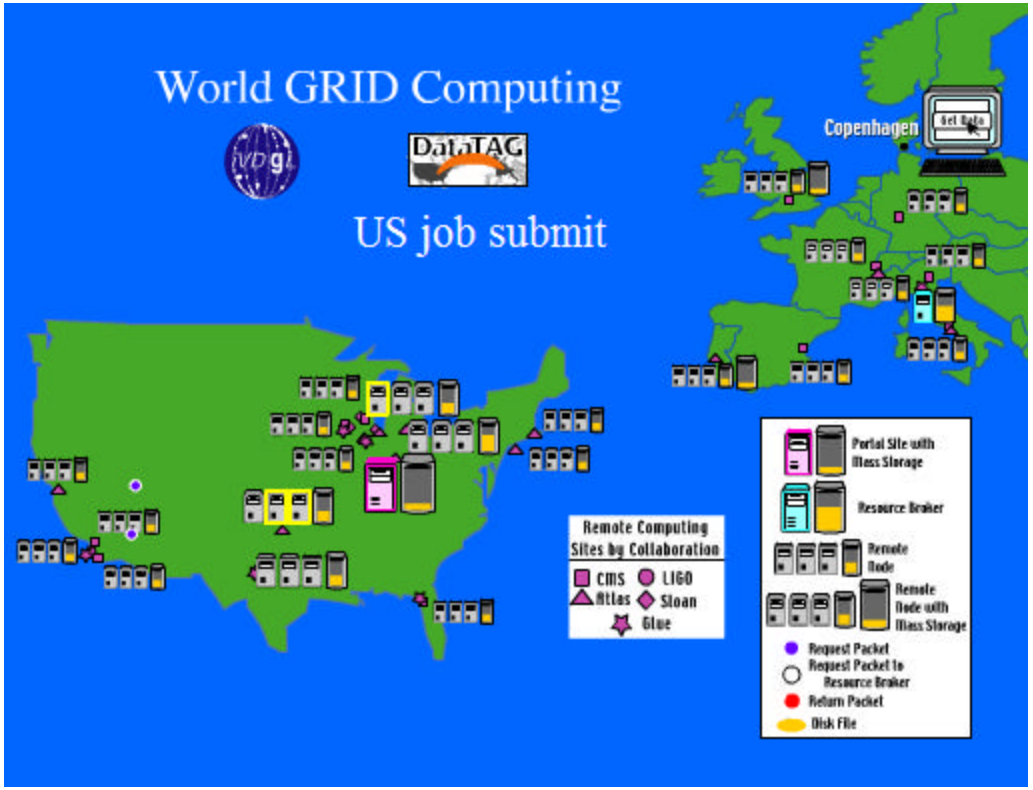

foster@mcs.anl.gov

NEESgrid Earthquake Engineering Collaboratory

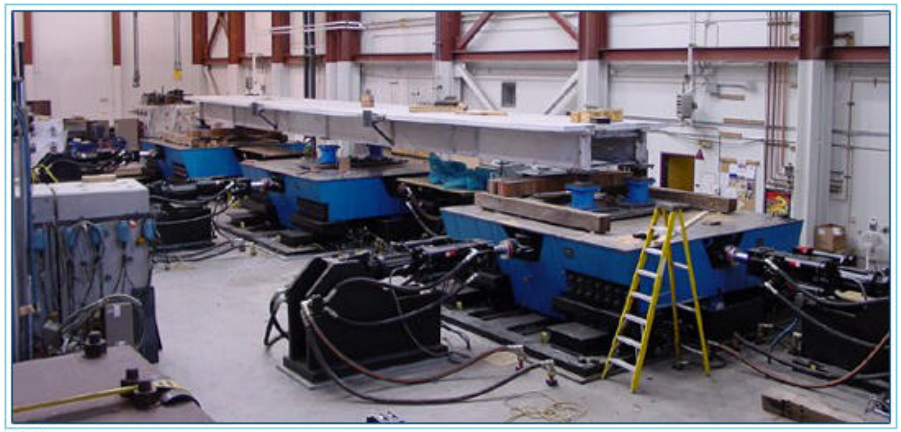

U.Nevada Reno

www.neesgrid.org

Remote Users
(Faculty, Students, Practitioners)

Instrumented Structures and Sites

Laboratory Equipment

High-Performance Network(s)

Field Equipment

Curated Data Repository

Global Connections
(fully developed FY 2005 – FY 2014)

Laboratory Equipment
(Faculty and Students)

Remote Users:
(K-12 Faculty and Students)

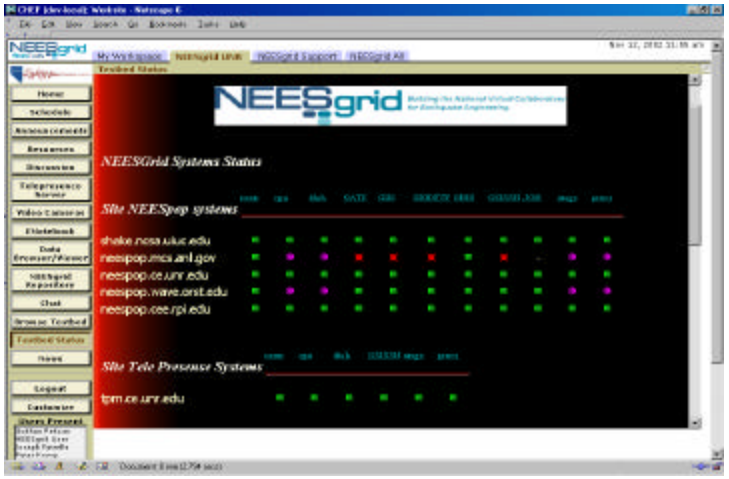

Grid Computing

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.

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

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.

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."

Grid Computing
By M. Mitchell Waldrop
May 2002

AN MIT ENTERPRISE
TECHNOLOGY REVIEW
Emerging Technologies and Their Impact

Hook enough computers together and what do you get? A new kind of utility that offers supercomputer processing on tap.

Is Internet history about to repeat itself?

Industrial Perspective on Grids: A Wide Range of Applications

Unique by Industry with Common Characteristics

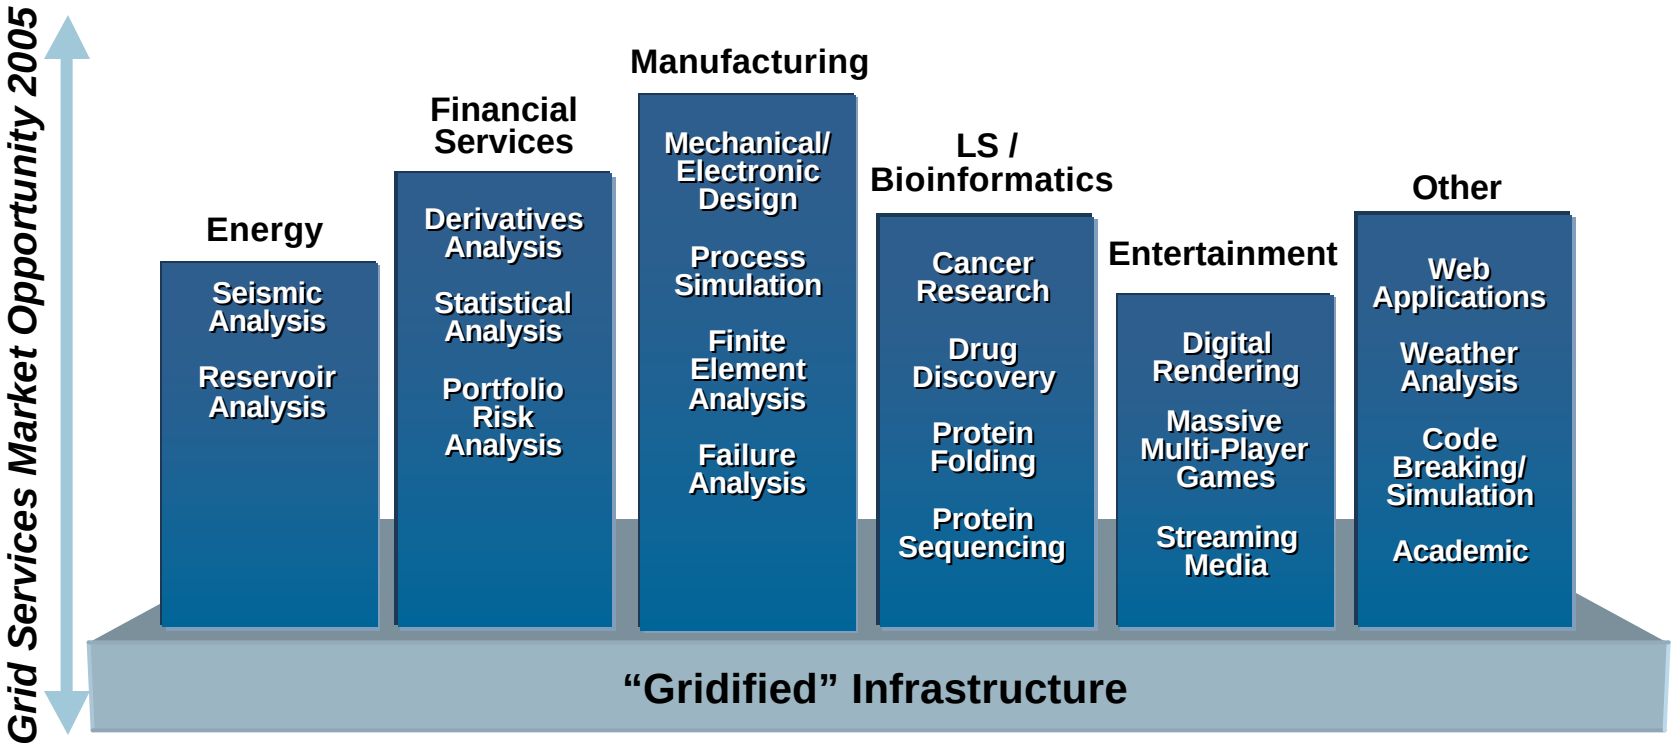

Sources: IDC, 2000 and Bear Stearns- Internet 3.0 - 5/01 Analysis by SAI

the globus project™
www.globus.org

Overview

- Grid: why and what
- **Evolution of Grid technology**
  - **Open Grid Service Architecture**
- Future directions
  - Towards lightweight VOs: dynamic trust relationships
  - Towards global knowledge communities: virtual data and dynamic workspaces

the globus project™
www.globus.org

Grids and Open Standards

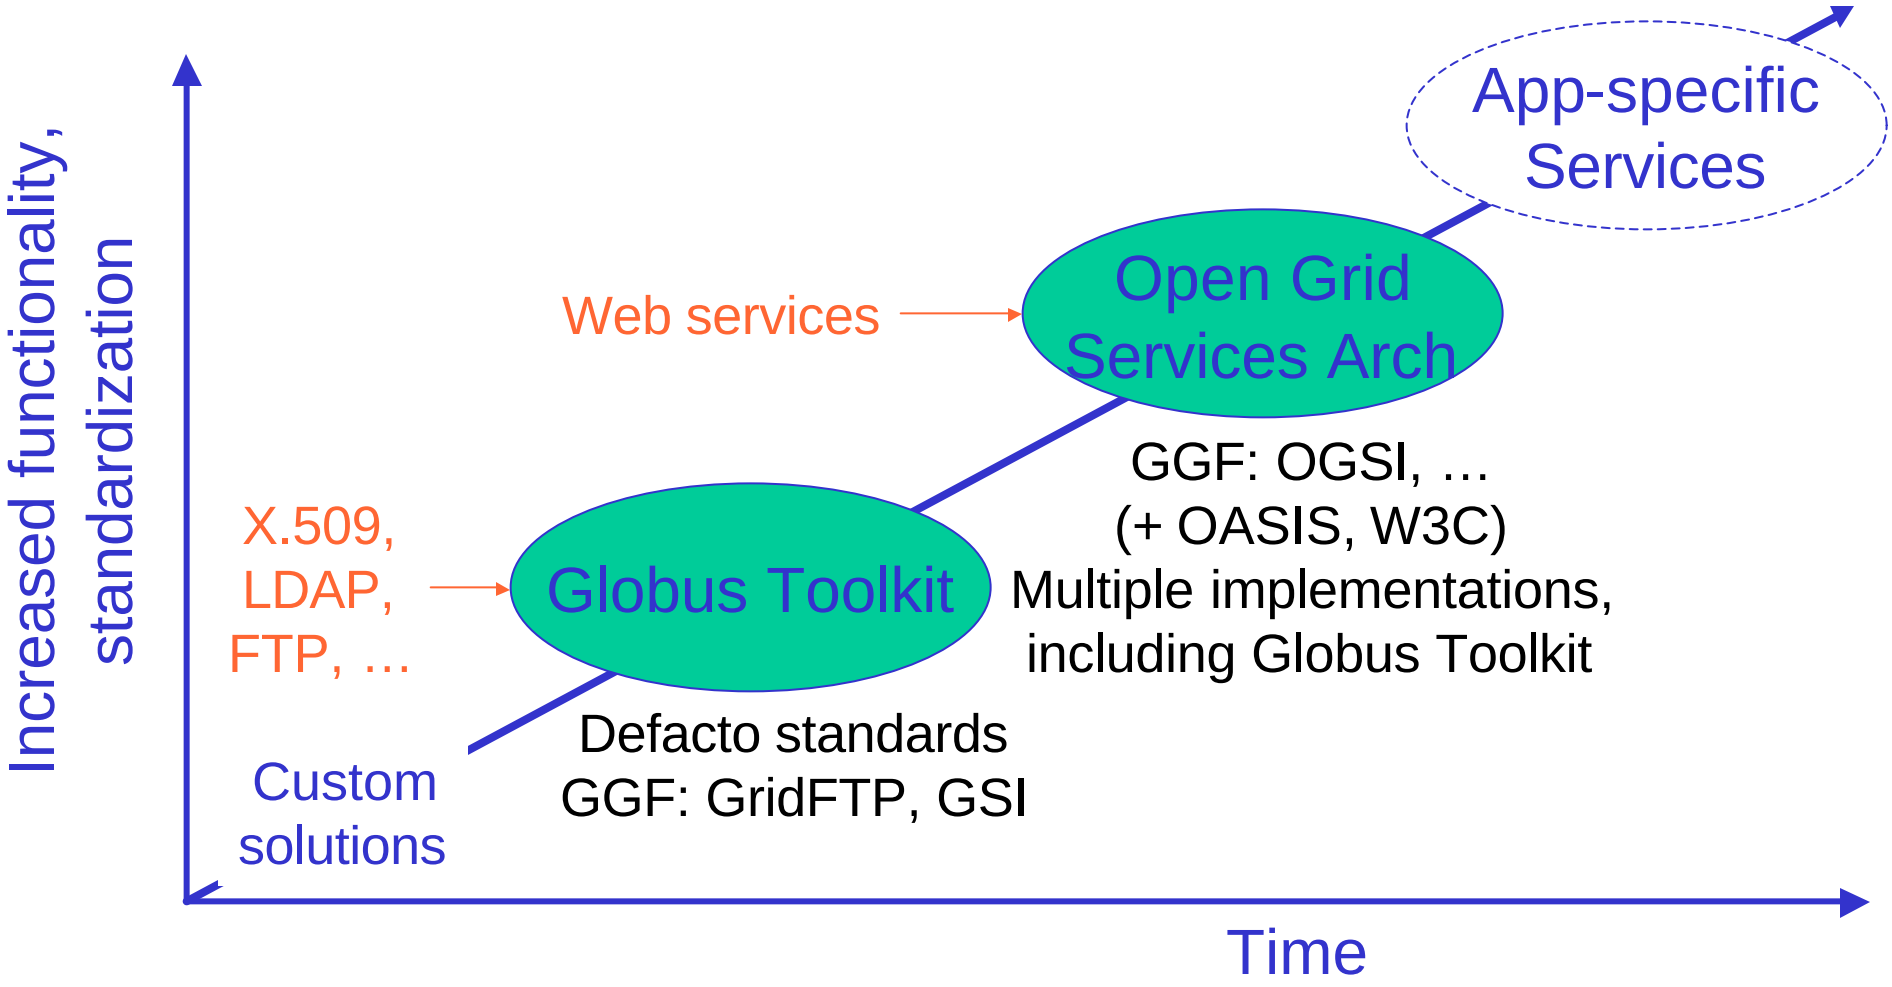

the globus project™
www.globus.org

Open Grid Services Architecture

- Service-oriented architecture
  - Key to virtualization, discovery, composition, local-remote transparency
- Leverage industry standards
  - Internet, Web services
- Distributed service management
  - A “component model for Web services”
- A framework for the definition of composable, interoperable services

“The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration”, Foster, Kesselman, Nick, Tuecke, 2002

the globus project™
www.globus.org

Web Services

- XML-based distributed computing technology
- Web service = a server process that exposes typed ports to the network
- Described by the Web Services Definition Language, an XML document that contains
  - Type of message(s) the service understands & types of responses & exceptions it returns
  - "Methods" bound together as "port types"
  - Port types bound to protocols as "ports"
- A WSDL document completely defines a service and how to access it

the globus project™
www.globus.org

OGSA Structure

- A standard substrate: the Grid service
  - Standard interfaces and behaviors that address key distributed system issues
  - A refactoring and extension of the Globus Toolkit protocol suite
- ... supports standard service specifications
  - Resource management, databases, workflow, security, diagnostics, etc., etc.
  - Target of current & planned GGF efforts
- ... and arbitrary application-specific services based on these & other definitions

Open Grid Services Infrastructure

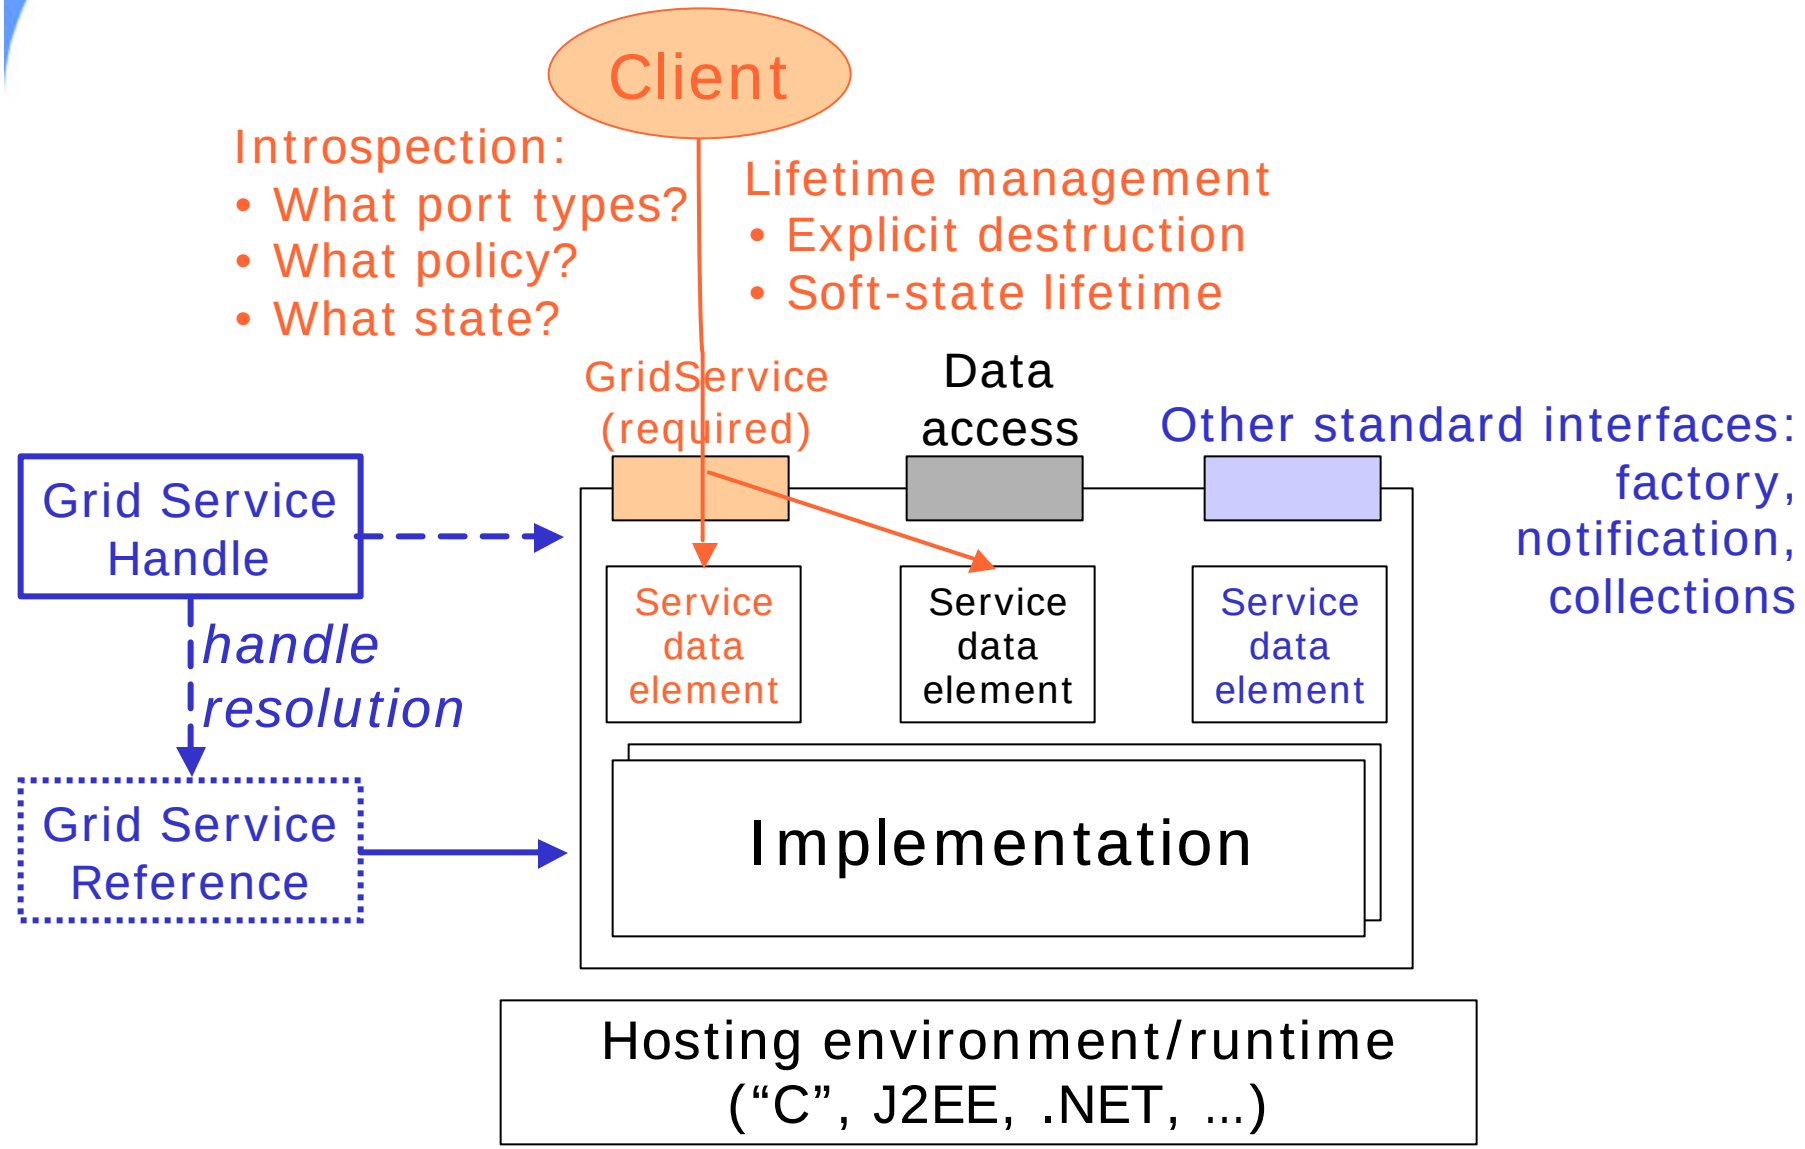

the globus project™
www.globus.org

Open Grid Services Infrastructure

GWD-R (draft-ggf-ogsi- gridservice-23)
Open Grid Services Infrastructure (OGSI)
http://www.ggf.org/ogsi-wg

Editors:
S. Tuecke, ANL
K. Czajkowski, USC/ISI
I. Foster, ANL
J. Frey, IBM
S. Graham, IBM
C. Kesselman, USC/ISI
D. Snelling, Fujitsu Labs
P. Vanderbilt, NASA
February 17, 2003

Open Grid Services Infrastructure (OGSI)

the globus project™
www.globus.org

Example: Reliable File Transfer Service

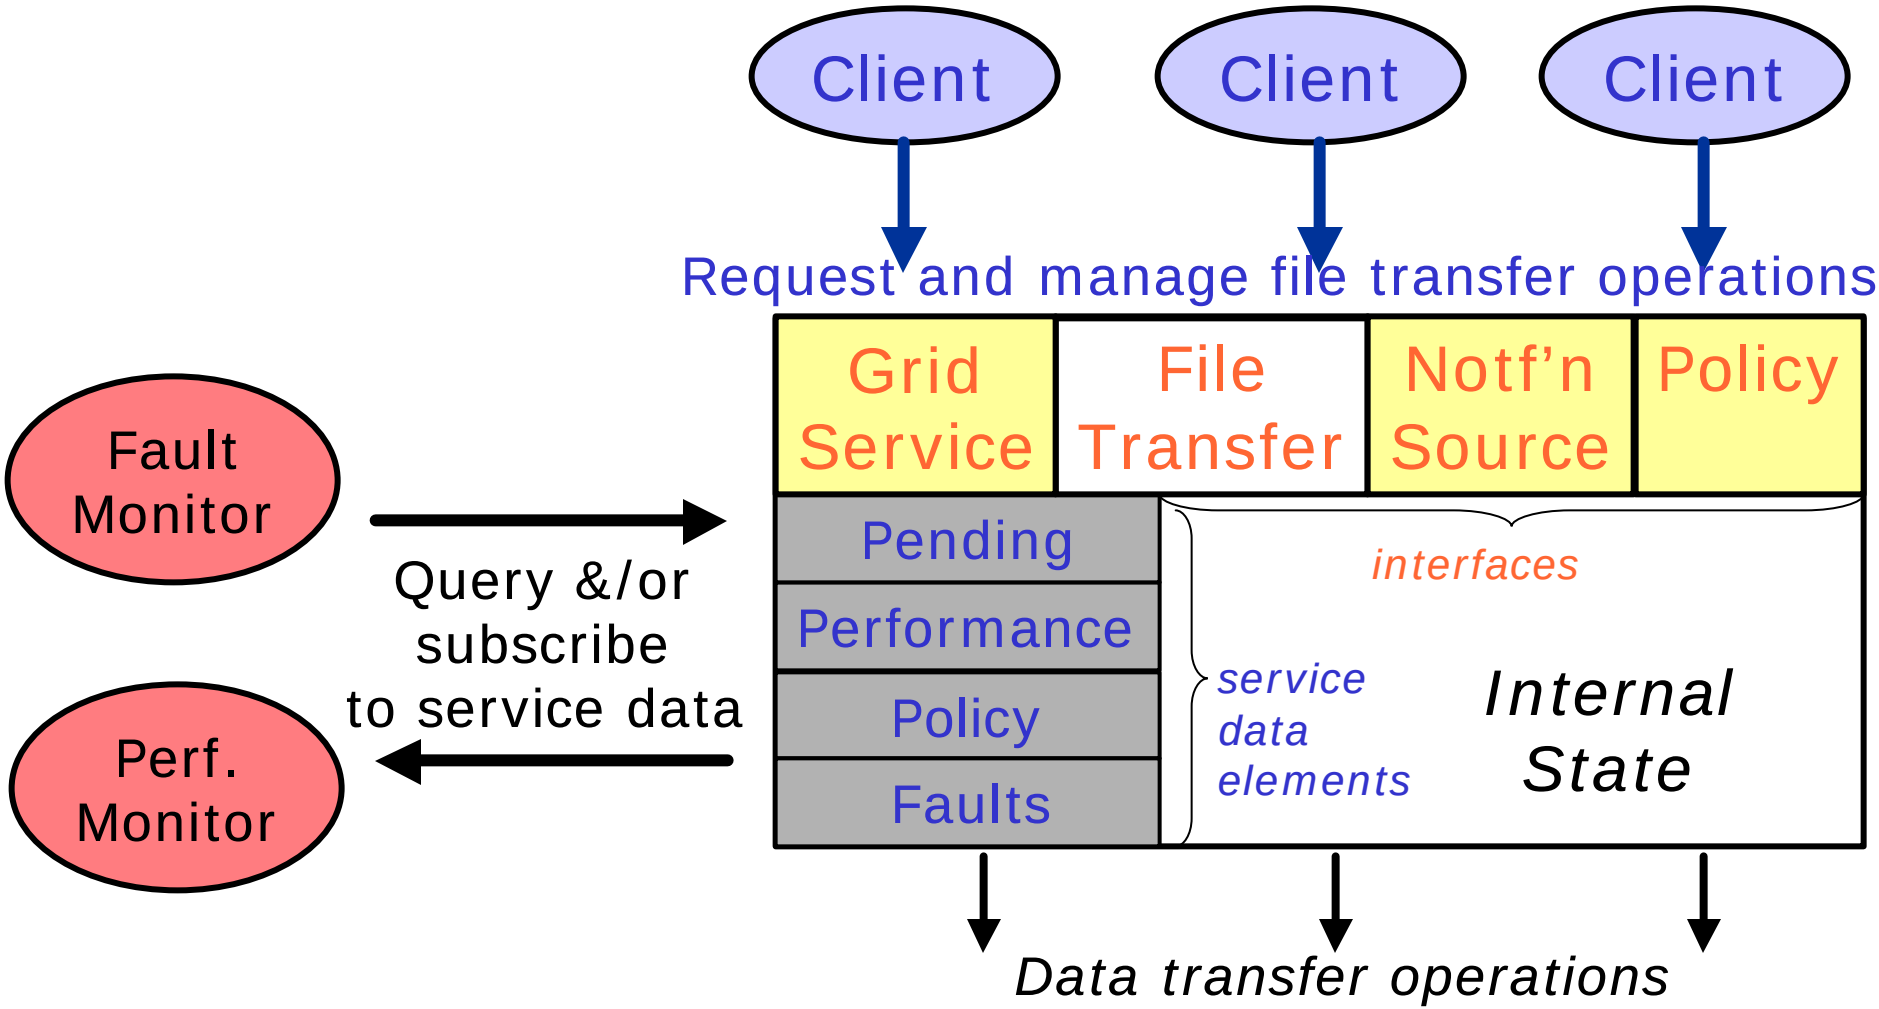

Open Grid Service Architecture: Next Steps

- ✓ Technical specifications
  - Open Grid Services Infrastructure is complete
  - Security, data access, Java binding, common resource models, etc., etc., in the pipeline
- ✓ Implementations and compliant products
  - Here: OGSA-based Globus Toolkit v3, …
  - Announced: IBM, Avaki, Platform, Sun, NEC, HP, Oracle, UD, Entropia, Insors, …, …
- ⌛ Rich set of service defns & implementations

Globus Toolkit v3 (GT3) Open Source OGSA Technology

- Implements OGSI interfaces
- Supports primary GT2 interfaces
  - High degree of backward compatibility
- Multiple platforms & hosting environments
  - J2EE, Java, C, .NET, Python
- New services
  - SLA negotiation, service registry, community authorization, data management, ...
- Rapidly growing adoption and contributions: "Linux for the Grid"

the globus project™
www.globus.org

Overview

- Grid: why and what
- Evolution of Grid technology
  - Open Grid Service Architecture
- **Future directions**
  - **Towards lightweight VOs: dynamic trust relationships**
  - **Towards global knowledge communities: virtual data and dynamic workspaces**

the globus project™
www.globus.org

Future Directions

- Grids are about computers, certainly
  - "On-demand" access to computing, etc.
  - Challenging future issues here: e.g., scale

the globus project™
www.globus.org

CMS Event Simulation Production

- Production Run on the Integration Testbed
  - Simulate 1.5 million full CMS events for physics studies: ~500 sec per event on 850 MHz processor
  - 2 months continuous running across 5 testbed sites
  - Managed by a single person at the US-CMS Tier 1

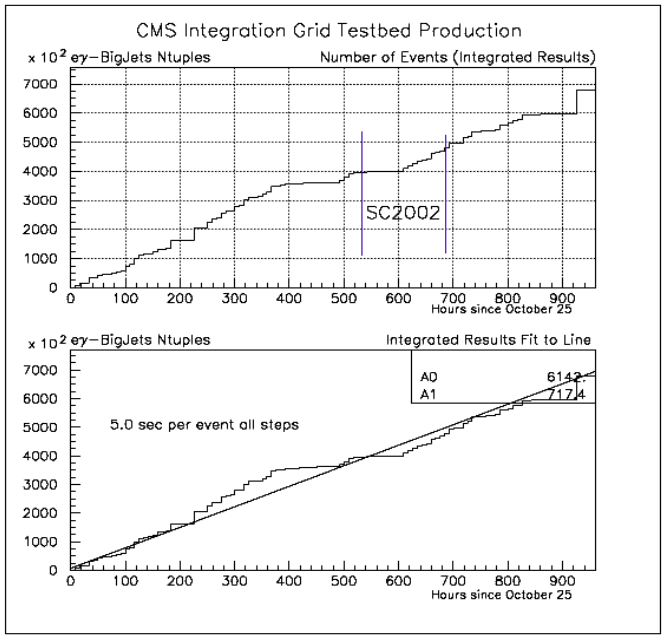

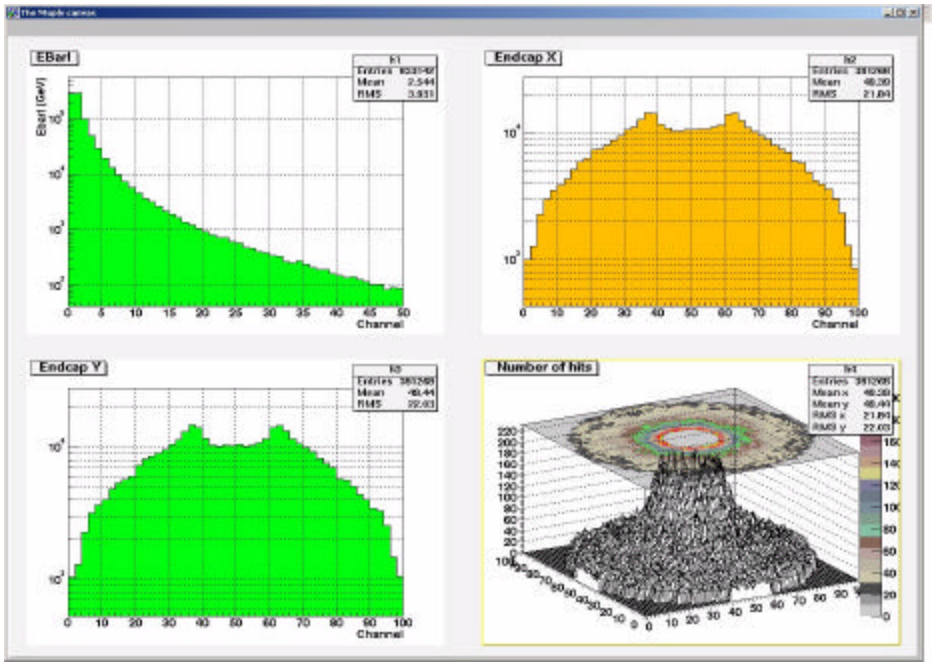

the globus project™
www.globus.org

CMS Event Simulation P… …on

- Production Run on the Integrati…
  - Simulate 1.5 million full CM… studies: ~500 sec per e…
  - 2 months continuou… …sites
  - Managed by a s… …S Tier 1

1.5 Million Events
Delivered to CMS Physicists!
(nearly 30 CPU years)

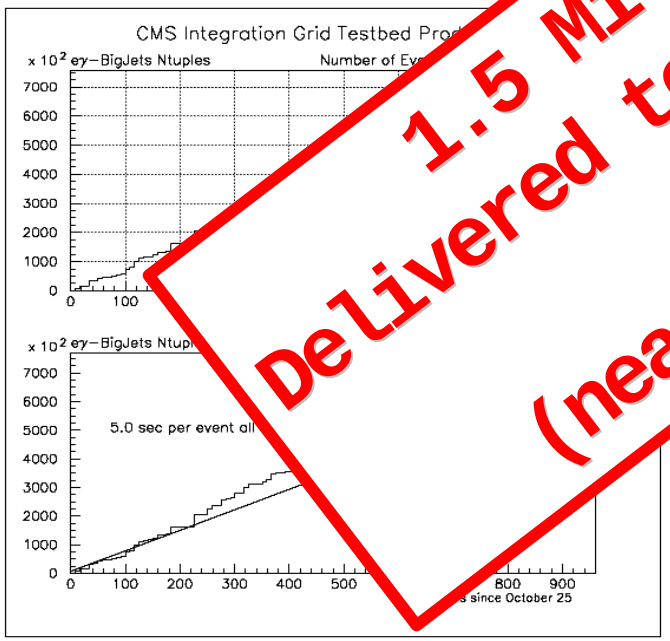

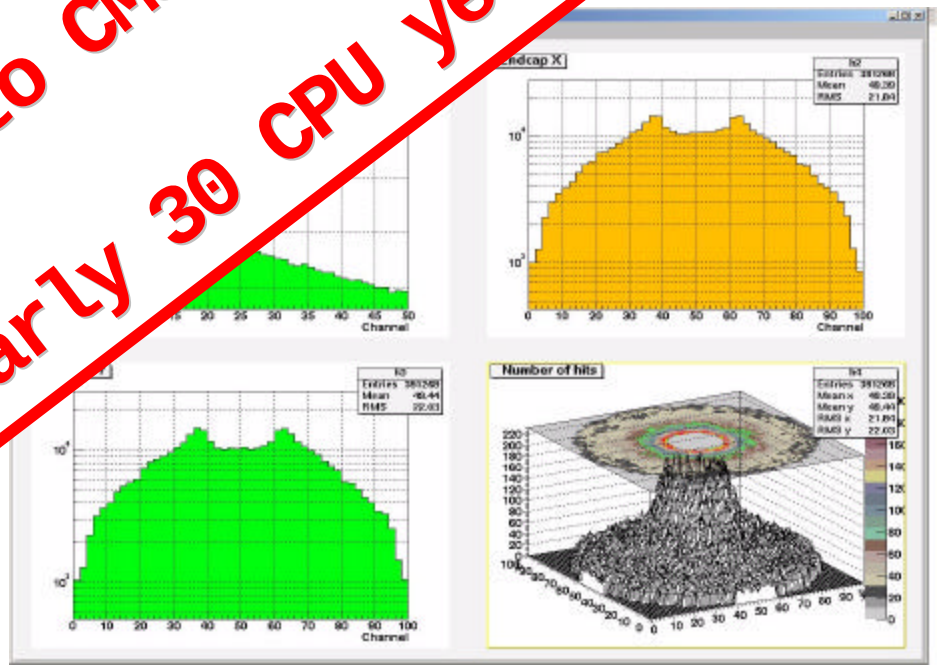

the globus project™
www.globus.org

Future Directions

- Grids are about computers, certainly
  - "On-demand" access to computing, etc.
  - Challenging future issues here: e.g., scale
- But they are ultimately about people, their activities, and their interactions
  - New interaction modalities supported by on-demand formation of lightweight VOs
  - New technologies needed: e.g., trust, security, data and knowledge integration
- Convergence of interest between "Access" and "Compute" Grids?

Global Knowledge Communities

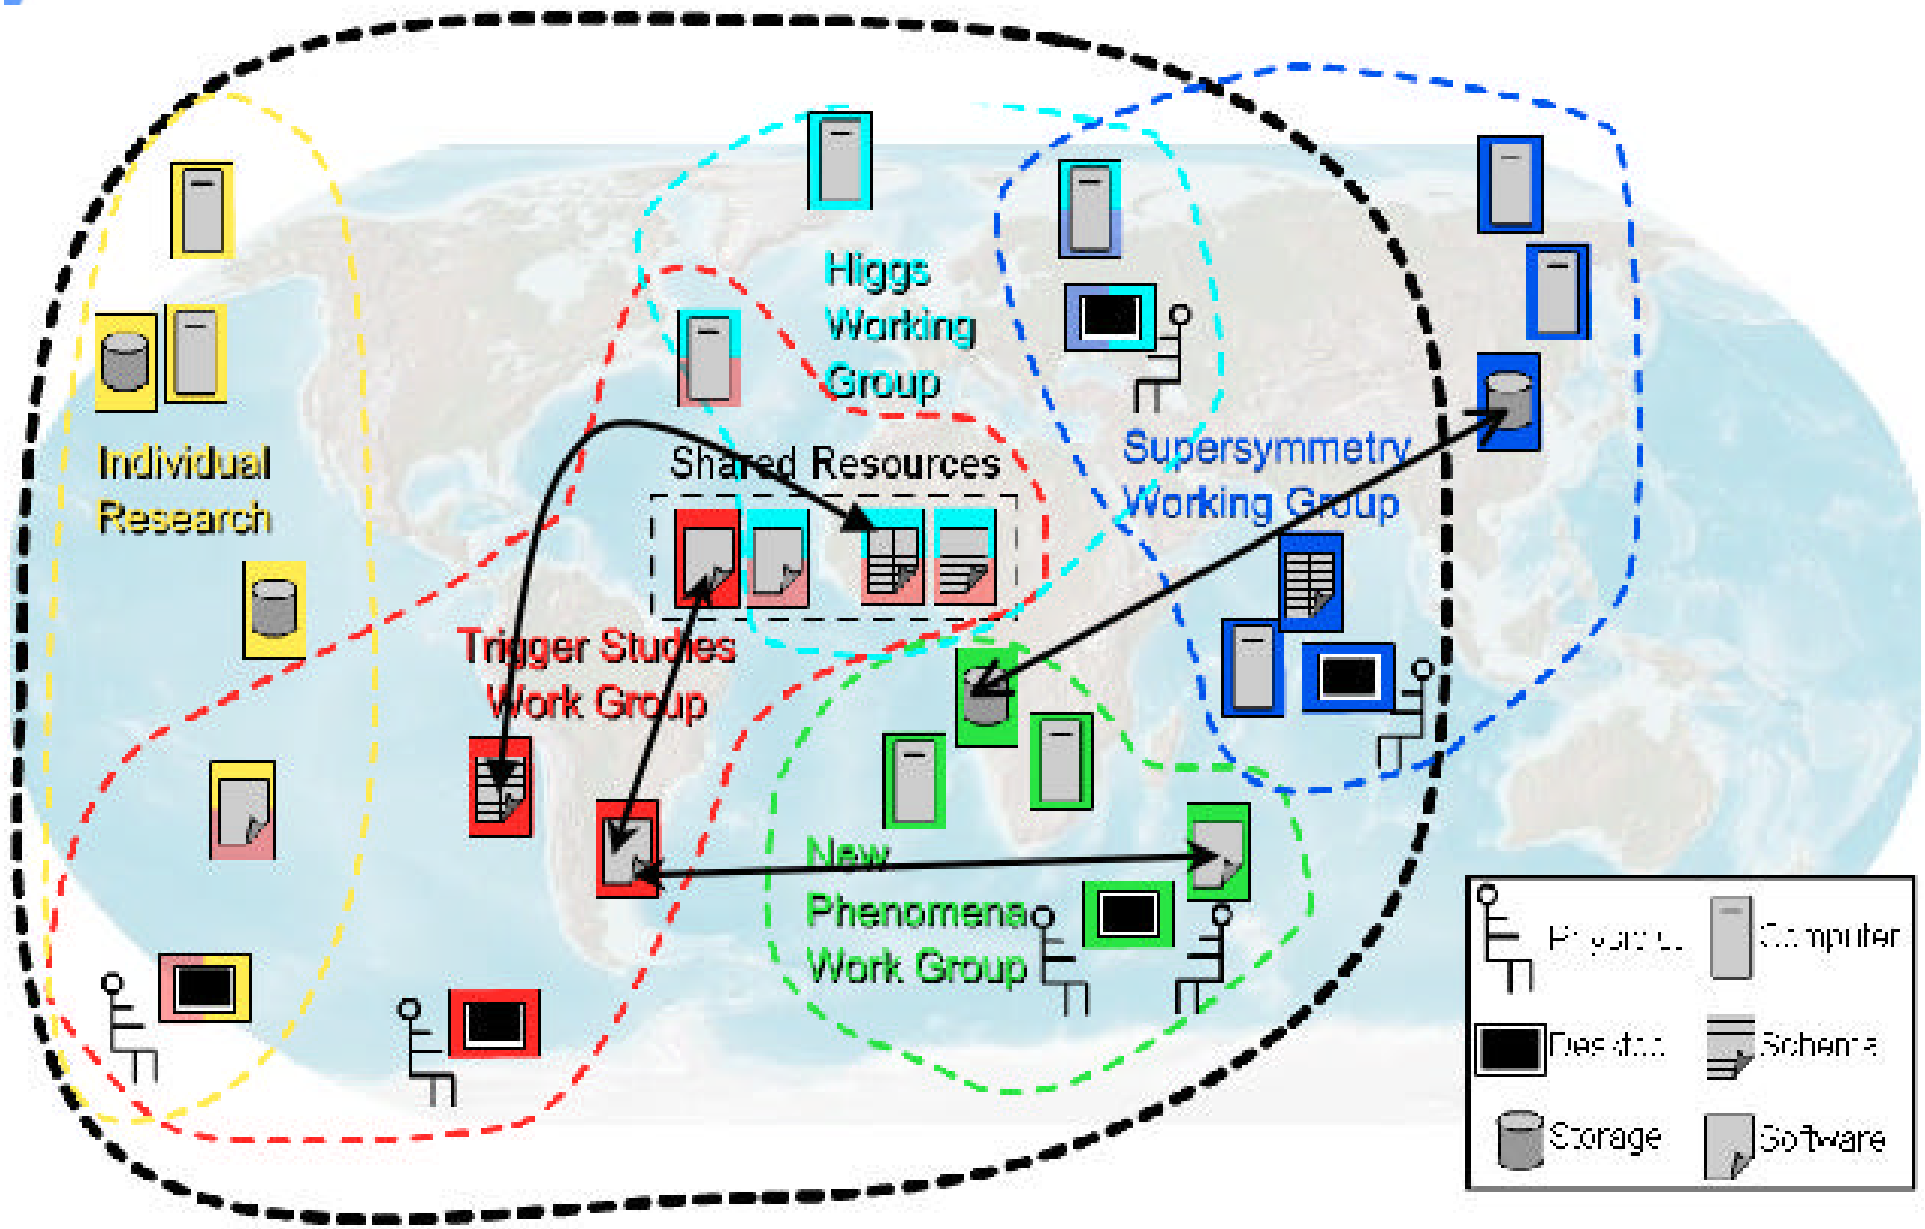

the globus project™
www.globus.org

Example Issue: Trust and Security

- Effective VO operation depends critically on
  - Trust: can I rely on you?
  - Protection mechanisms to govern actions
- Suffers from VO-organization policy mismatch
- Goal: collaborations no longer defined by slow centralized mechanisms but can
  - form spontaneously;
  - be managed in a distributed manner; and
  - be protected by an infrastructure that maintains and enforces trust relationships

Grid Security Services

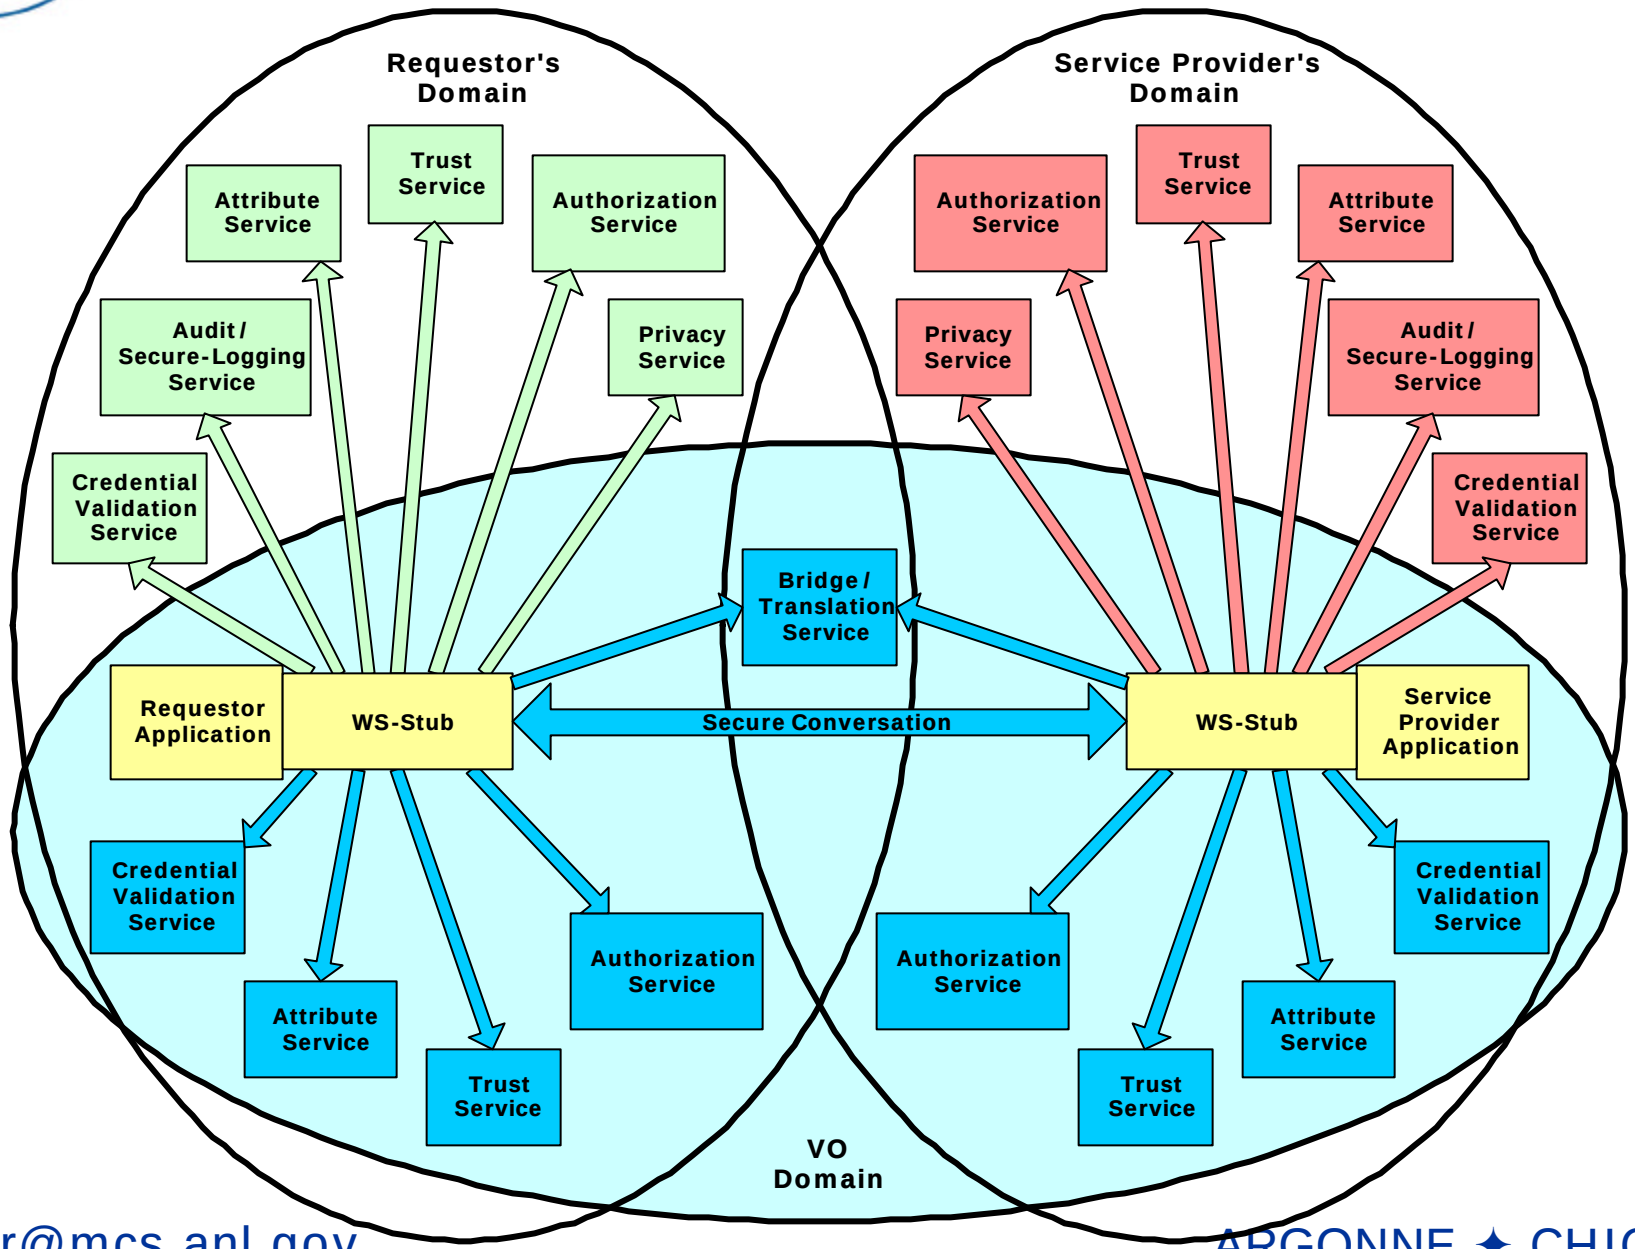

Understanding and Enhancing VO Trust and Security

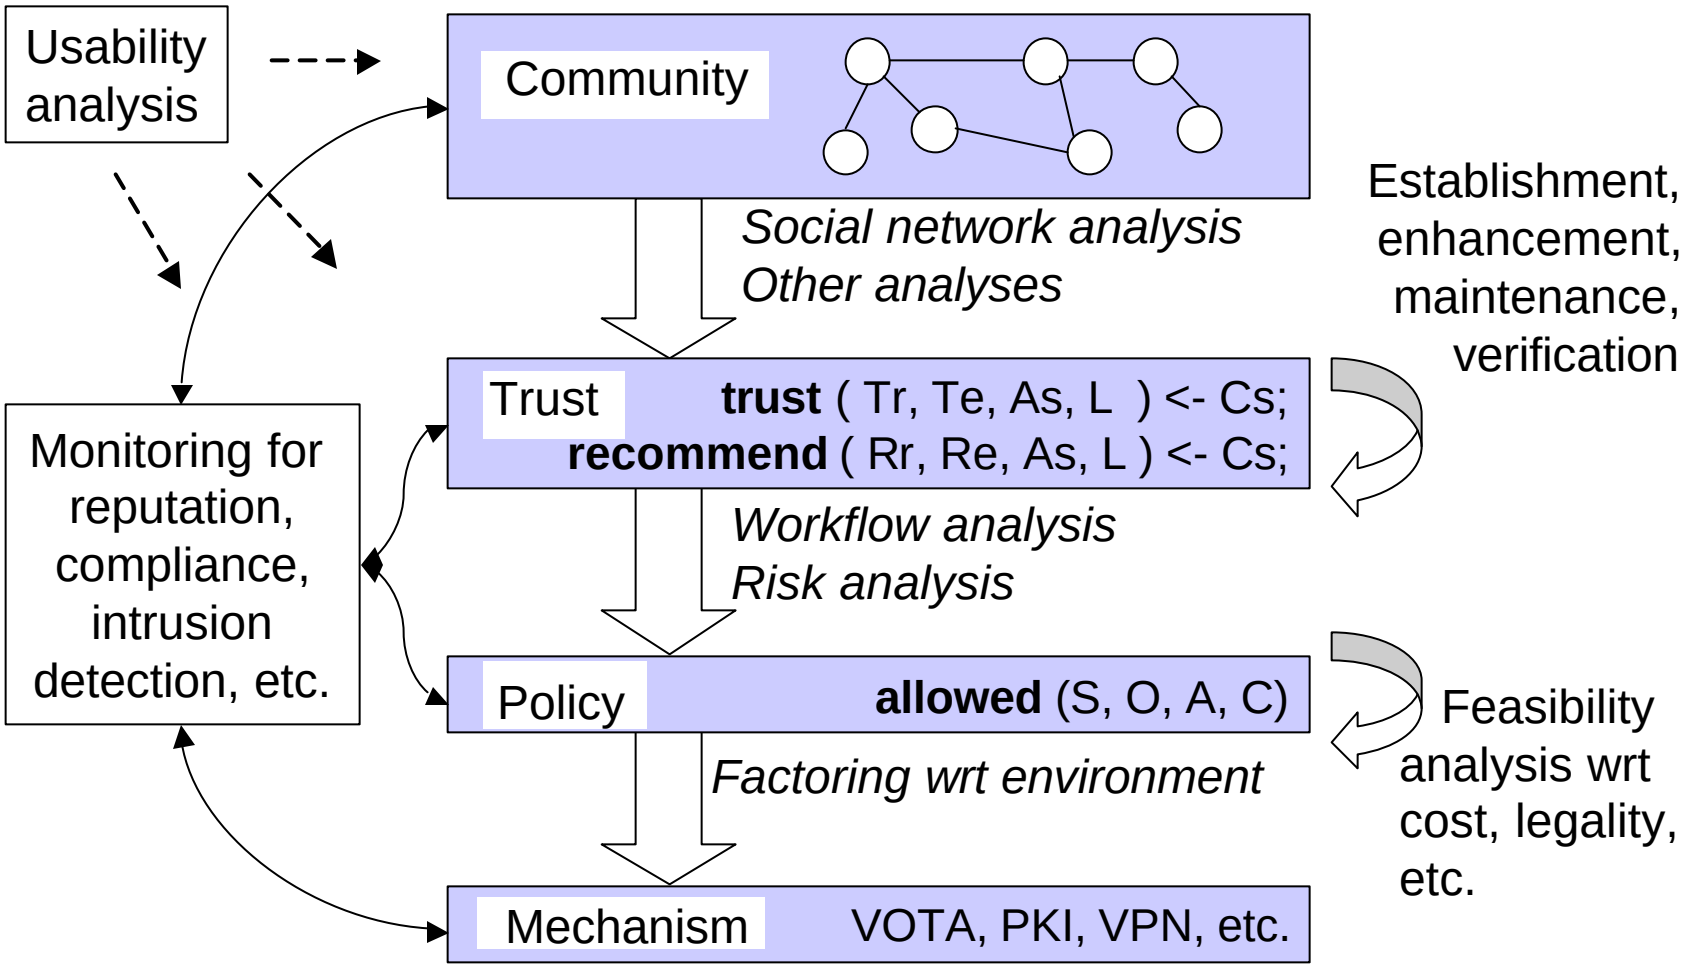

Example: Virtual Data for Collaborative Science

- Much collaboration is concerned with the development & use of knowledge, whether
  - Programs for data analysis and generation
  - Computations involving those programs
  - Metadata concerning data, programs, computations—and their interrelationships
- In a distributed, heterogeneous, fractal (?) environment with widely varying
  - Data and analysis program formats
  - Degrees of formality and scale
  - Scientific goals and sharing policies

Sloan Digital Sky Survey Production System

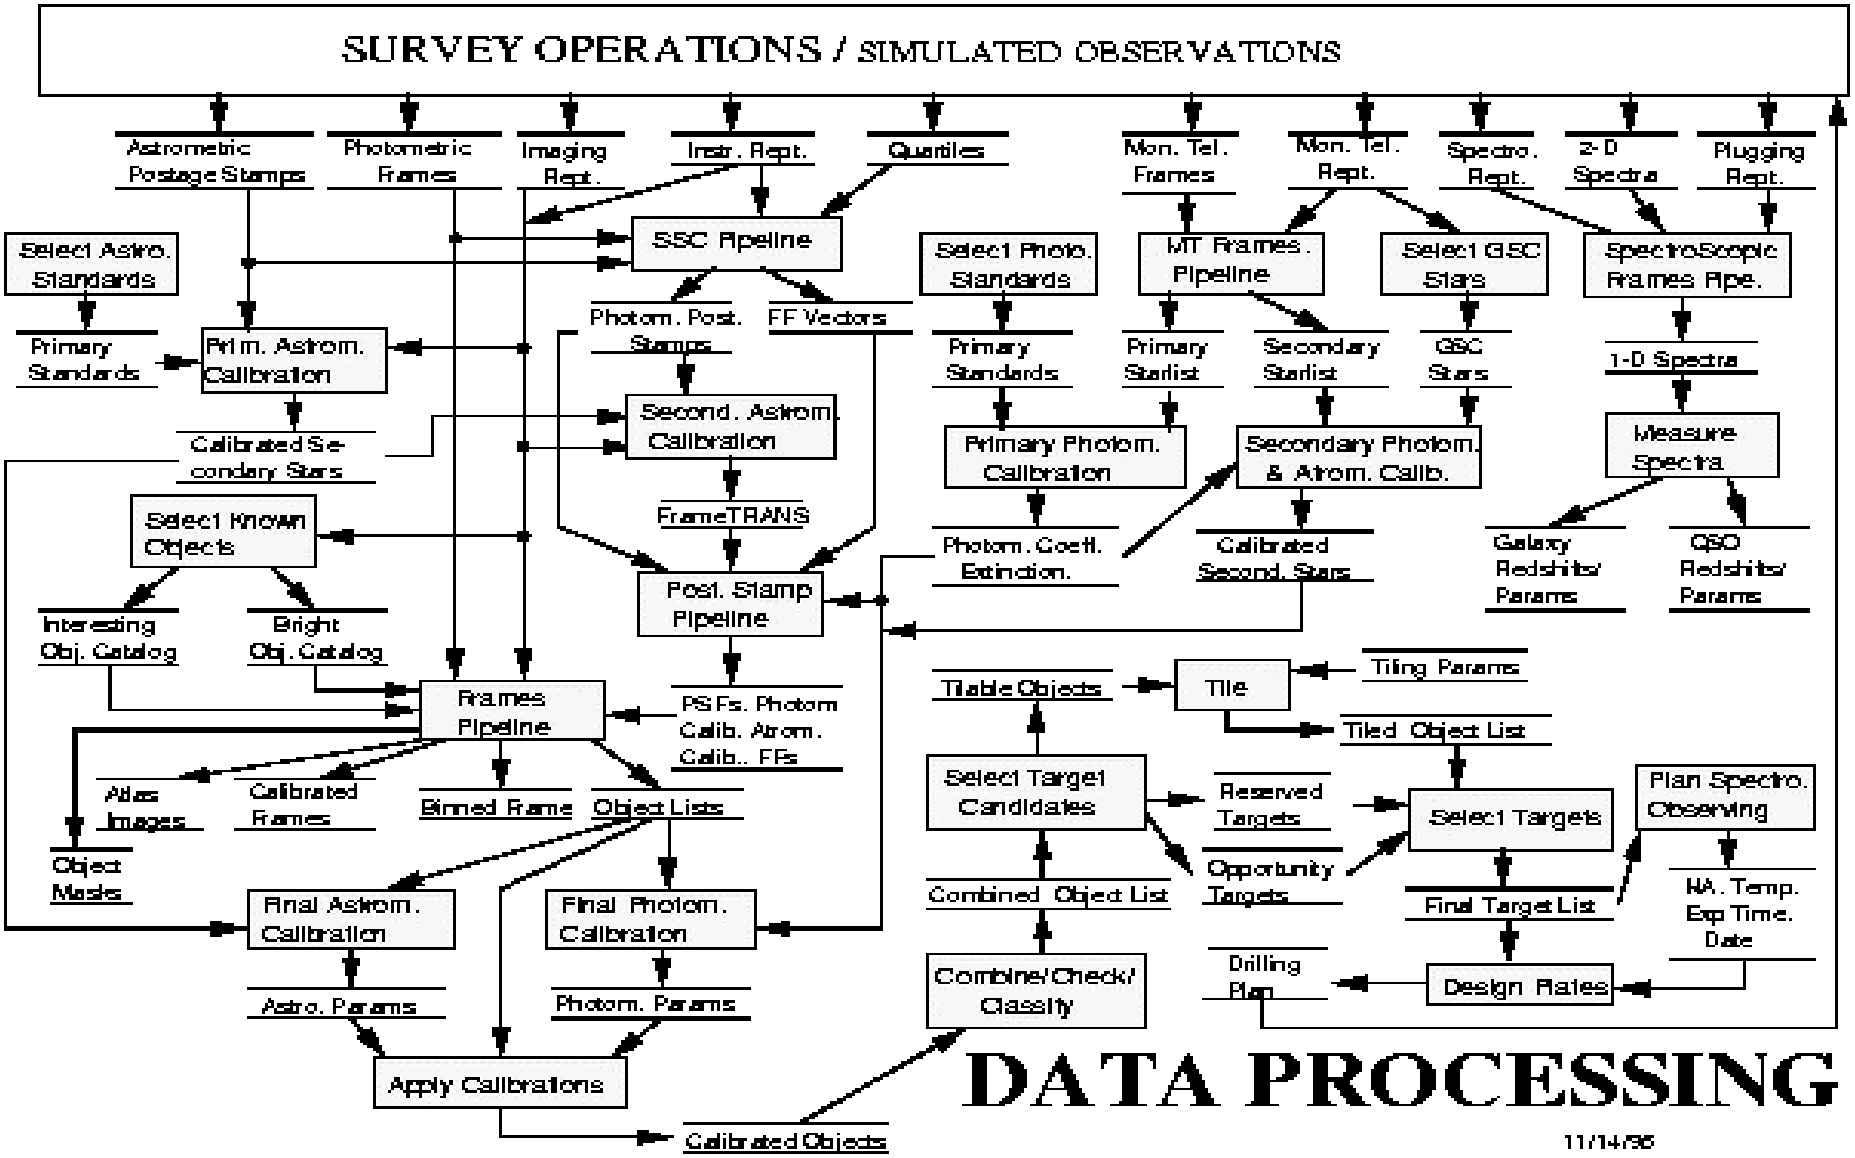

Virtual Data Concept

- Capture and manage information about relationships among
  - Data (of widely varying representations)
  - Programs (& their execution needs)
  - Computations (& execution environments)
- Apply this information to, e.g.
  - Discovery: Data and program discovery
  - Workflow: Structured paradigm for organizing, locating, specifying, & requesting data
  - Explanation: provenance
  - Planning and scheduling
  - Other uses we haven't thought of

Motivations

"I've come across some interesting data, but I need to understand the nature of the corrections applied when it was constructed before I can trust it for my purposes."

"I've detected a calibration error in an instrument and want to know which derived data to recompute."

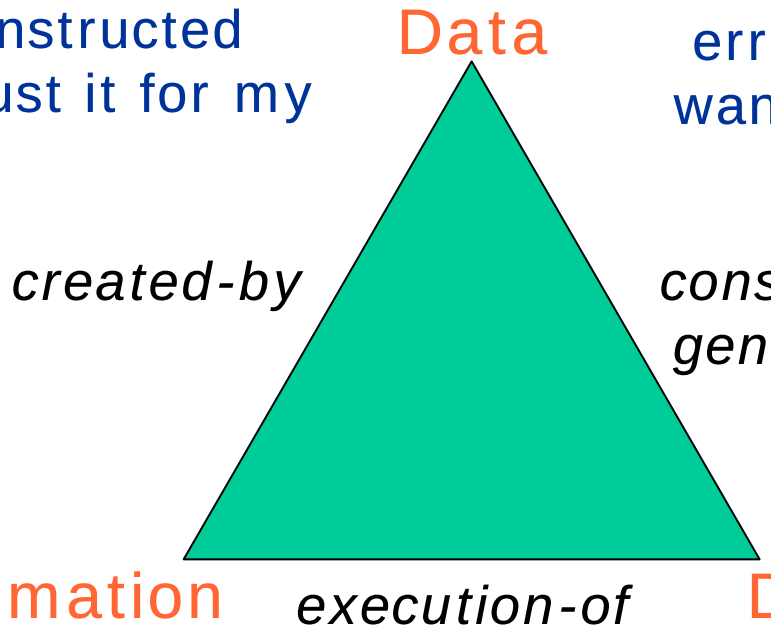

"I want to search an astronomical database for galaxies with certain characteristics. If a program that performs this analysis exists, I won't have to write one from scratch."

"I want to apply an astronomical analysis program to millions of objects. If the results already exist, I'll save weeks of computation."

Example: Sloan Galaxy Cluster Analysis

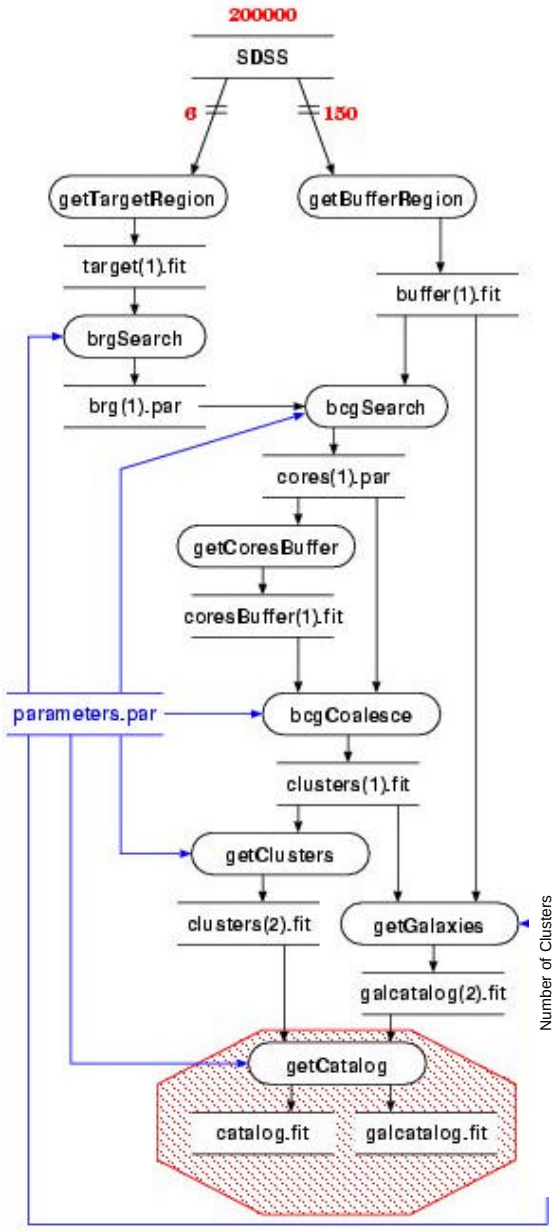

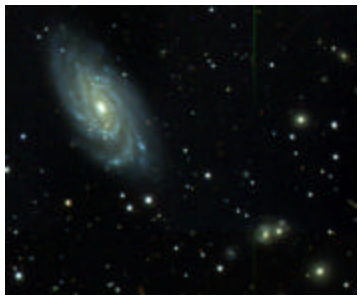

Sloan Data

DAG

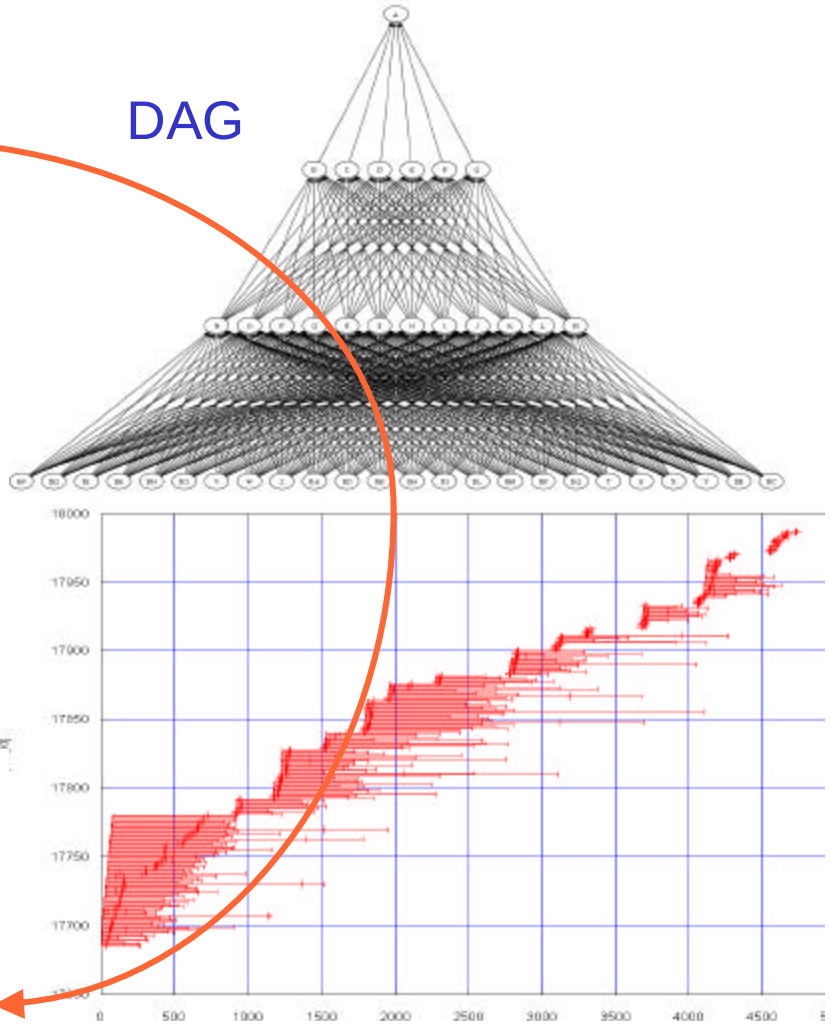

Galaxy cluster size distribution

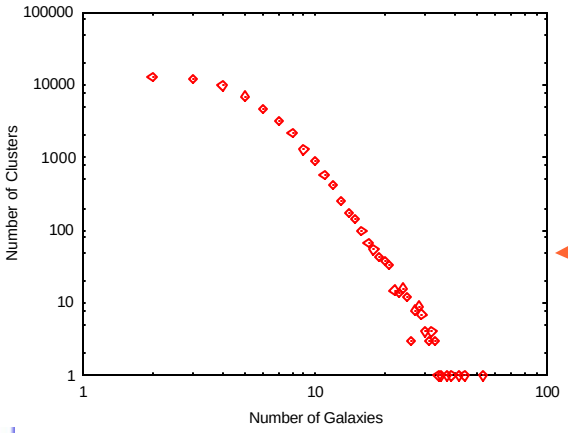

Jim Annis, Steve Kent, Vijay Sehkri, Fermilab; Michael Milligan, Yong Zhao, Chicago

the globus project™
www.globus.org

Integrating Provenance Data

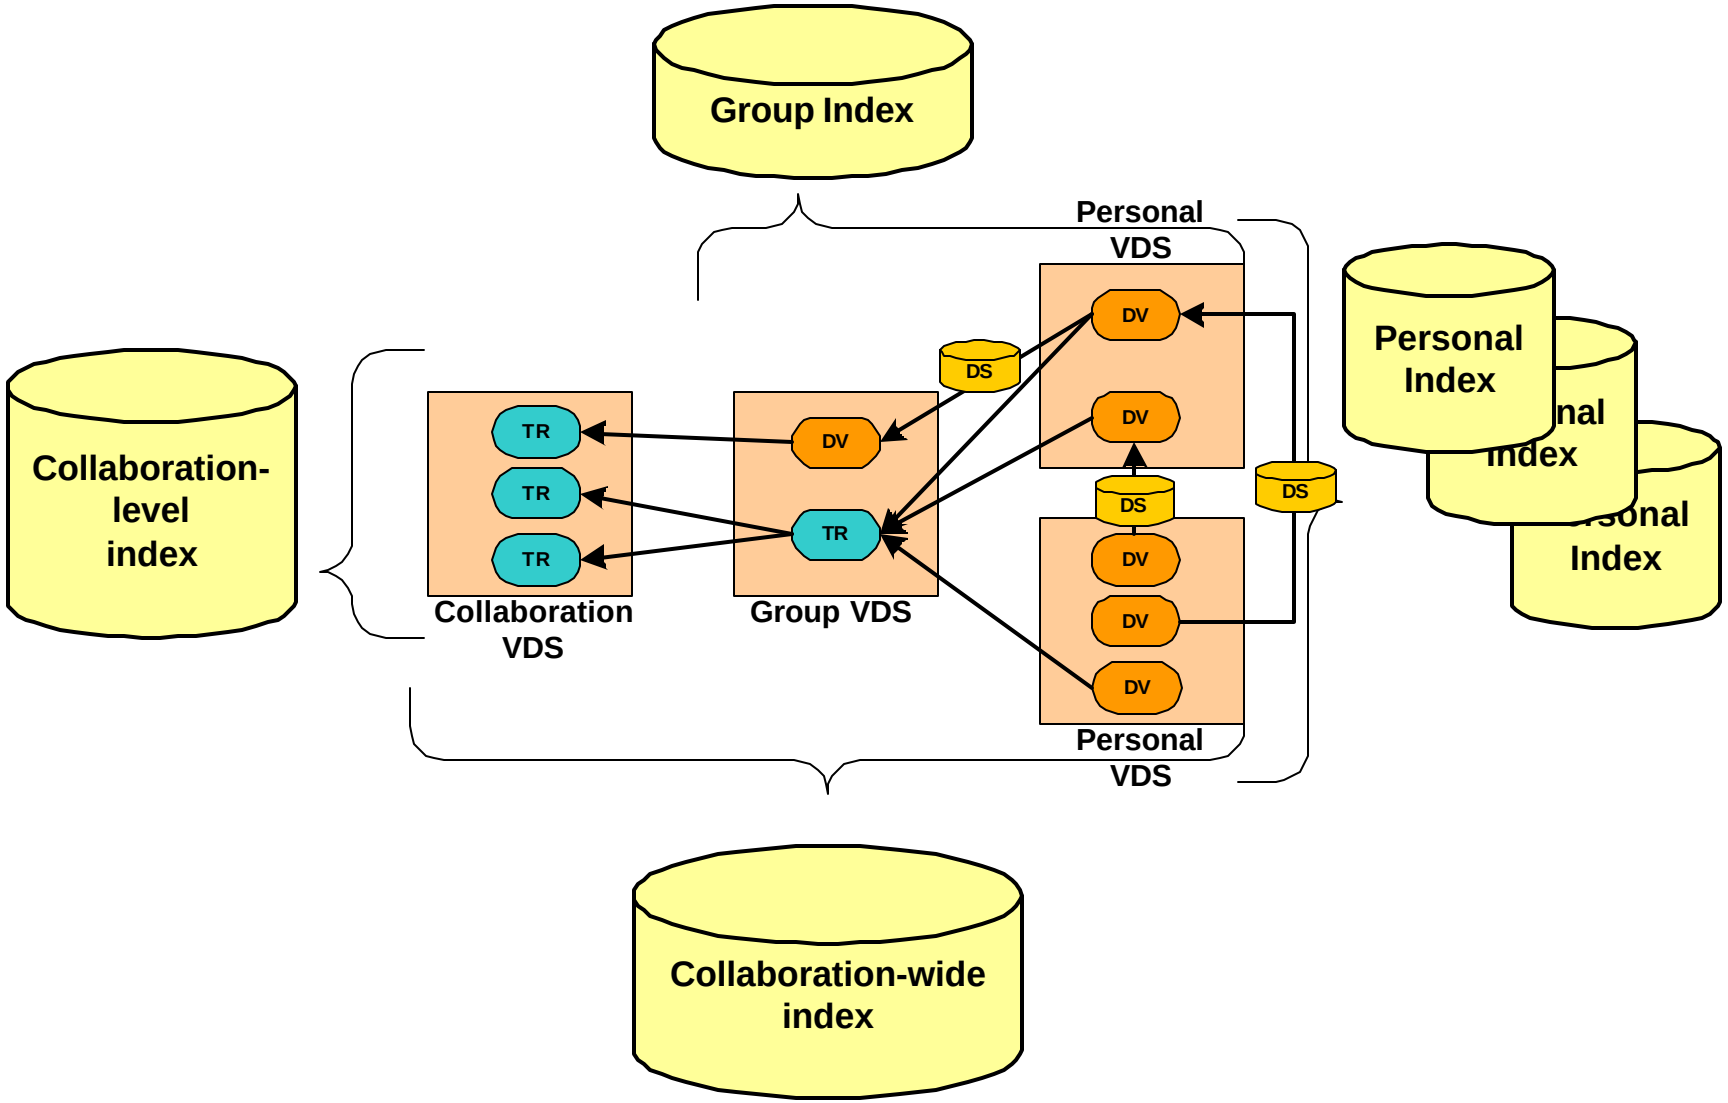

the globus project™
www.globus.org

Summary

- Yesterday
  - Small, static communities, primarily in science
  - Focus on sharing of computing resources
  - Globus Toolkit as technology base
- Today
  - Larger communities in science; early industry
  - Focused on sharing of data and computing
  - Open Grid Services Architecture emerging
- Tomorrow
  - Large, dynamic, diverse communities that share a wide variety of services, resources, data
  - New issues: Trust, distributed RM, knowledge