

Building an Open Grid: A Status Report

Ian Foster

Argonne National Lab
University of Chicago
Globus Alliance

www.mcs.anl.gov/~foster

CERN, September 9, 2003

Abstract

Grid technologies and infrastructure support the integration of services and resources within and among enterprises, and thus allow new approaches to problem solving and interaction within distributed, multi-organizational collaborations. Sustained effort by computer scientists and application developers has resulted in the creation of a substantial open source technology, numerous infrastructure deployments, a vibrant international community, and significant application success stories. Long-term success now depends critically on three issues: open standards, open software, and open infrastructure. I discuss current efforts and future directions in each area, referring in particular to recent developments in "cyber-infrastructure" in the U.S., EGEE in Europe, Open Grid Services Architecture standards, and adoption within scientific communities beyond the physical sciences.

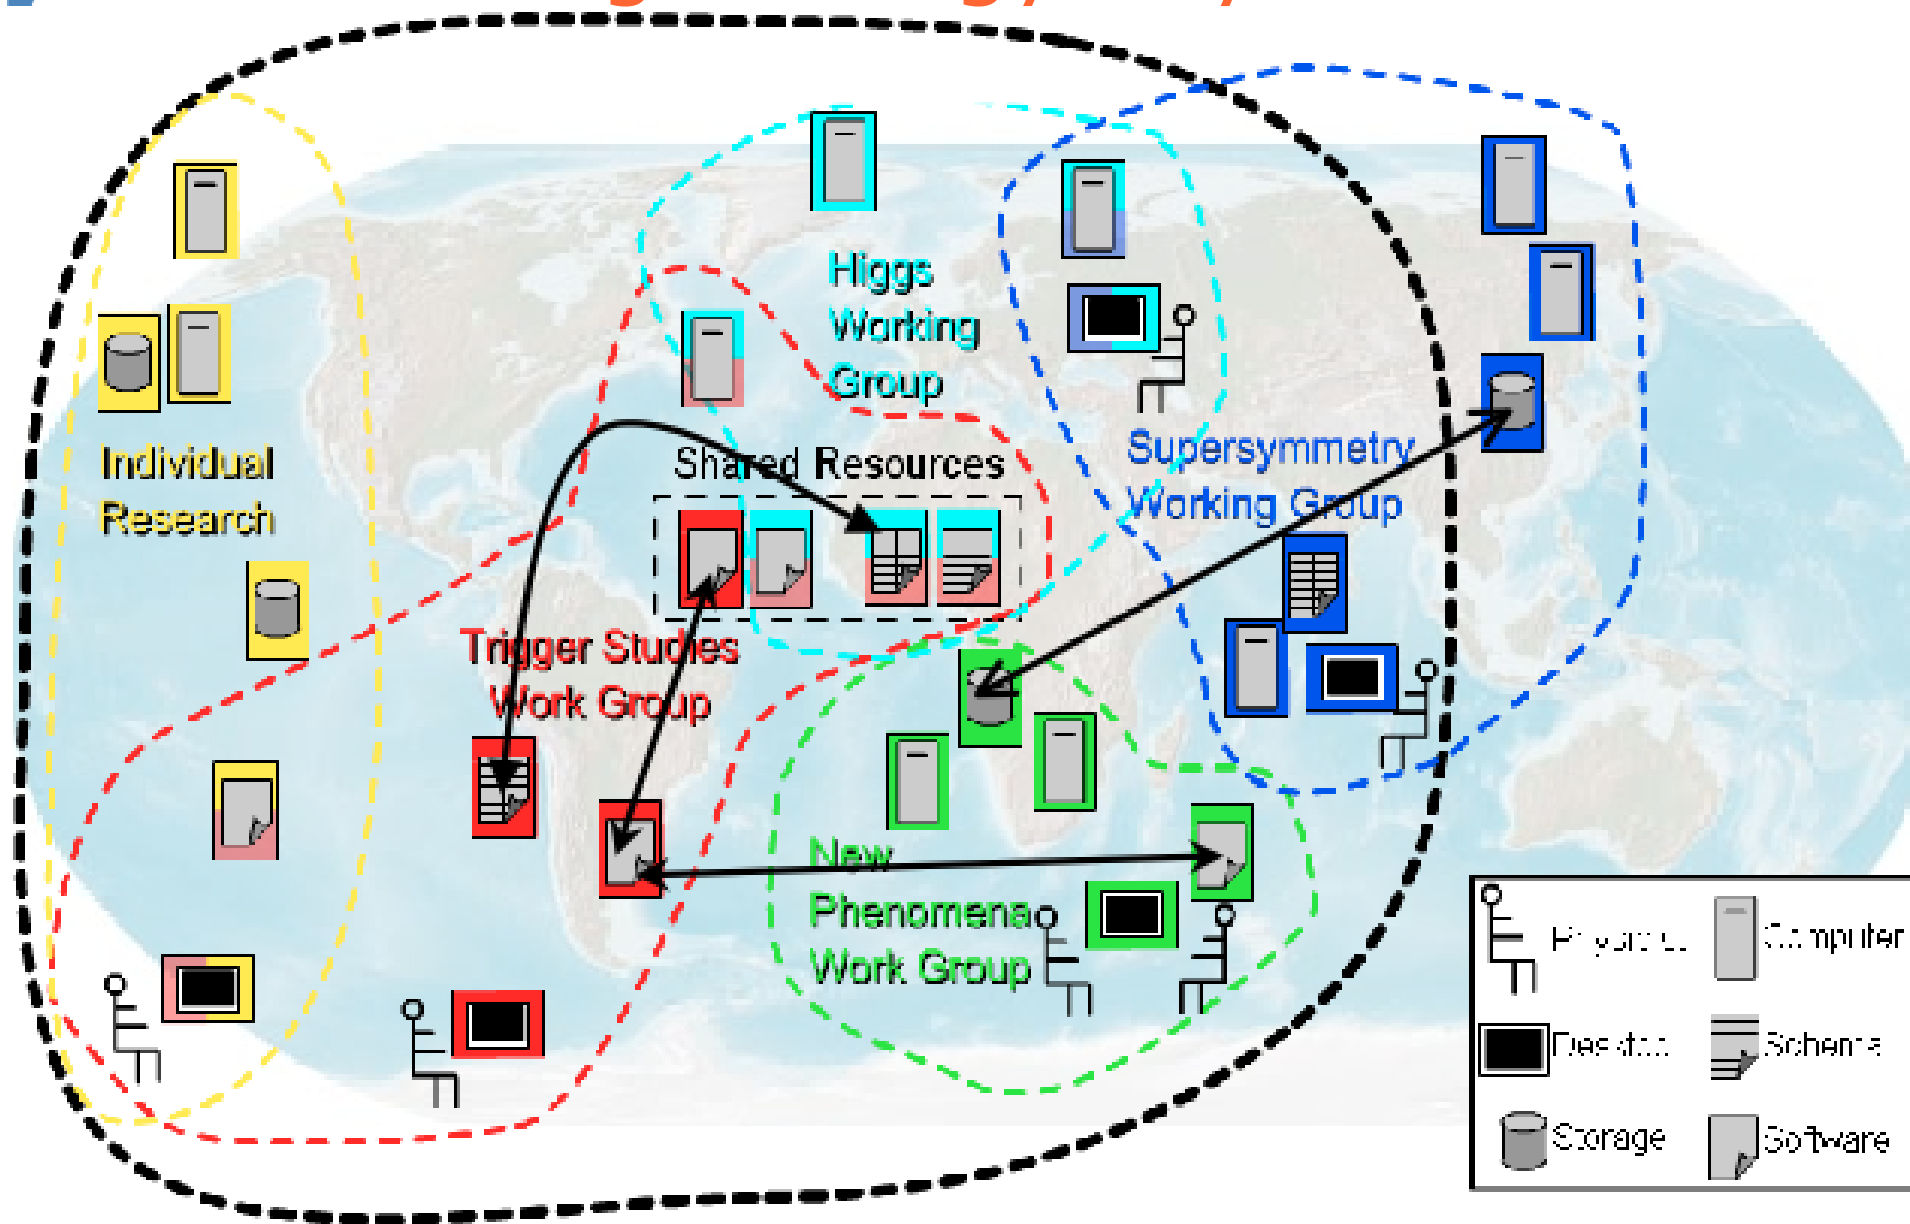
It's Easy to Forget How Different 2003 is From 1993

- Ubiquitous Internet: 100+ million hosts
 - ◆ Collaboration & resource sharing the norm
- Ultra-high-speed networks: 10+ Gb/s
 - ◆ Global optical networks
- Enormous quantities of data: Petabytes
 - ◆ For an increasing number of communities, gating step is not collection but analysis
- Huge quantities of computing: 100+ Top/s
 - ◆ Moore's law gives us all supercomputers

Consequence: The Emergence of Global Knowledge Communities

- Teams organized around common goals
 - ◆ Communities: “Virtual organizations”
- With diverse membership & capabilities
 - ◆ Heterogeneity is a strength not a weakness
- And geographic and political distribution
 - ◆ No location/organization possesses all required skills and resources
- Must adapt as a function of the situation
 - ◆ Adjust membership, reallocate responsibilities, renegotiate resources

For Example: High Energy Physics



Resource Integration as a Fundamental Challenge


Many sources
of data, services,
computation

Registries organize
services of interest
to a community

Security
service

Data integration activities
may require access to, &
exploration/analysis of, data
at many locations

foster@mcs.anl.gov




Security & policy
must underlie access
& management
decisions

Resource management
is needed to ensure
progress & arbitrate
competing demands

Policy
service


Exploration & analysis
may involve complex,
multi-step workflows

Grid Technologies Promise to Address Key Requirements




- Infrastructure (“middleware”) for establishing, managing, and evolving multi-organizational federations
 - ◆ Dynamic, autonomous, domain independent
 - ◆ On-demand, ubiquitous access to computing, data, and services
- Mechanisms for creating and managing workflow within such federations
 - ◆ New capabilities constructed dynamically and transparently from distributed services
 - ◆ Service-oriented, virtualization

Realizing the Promise: Building an Open Grid



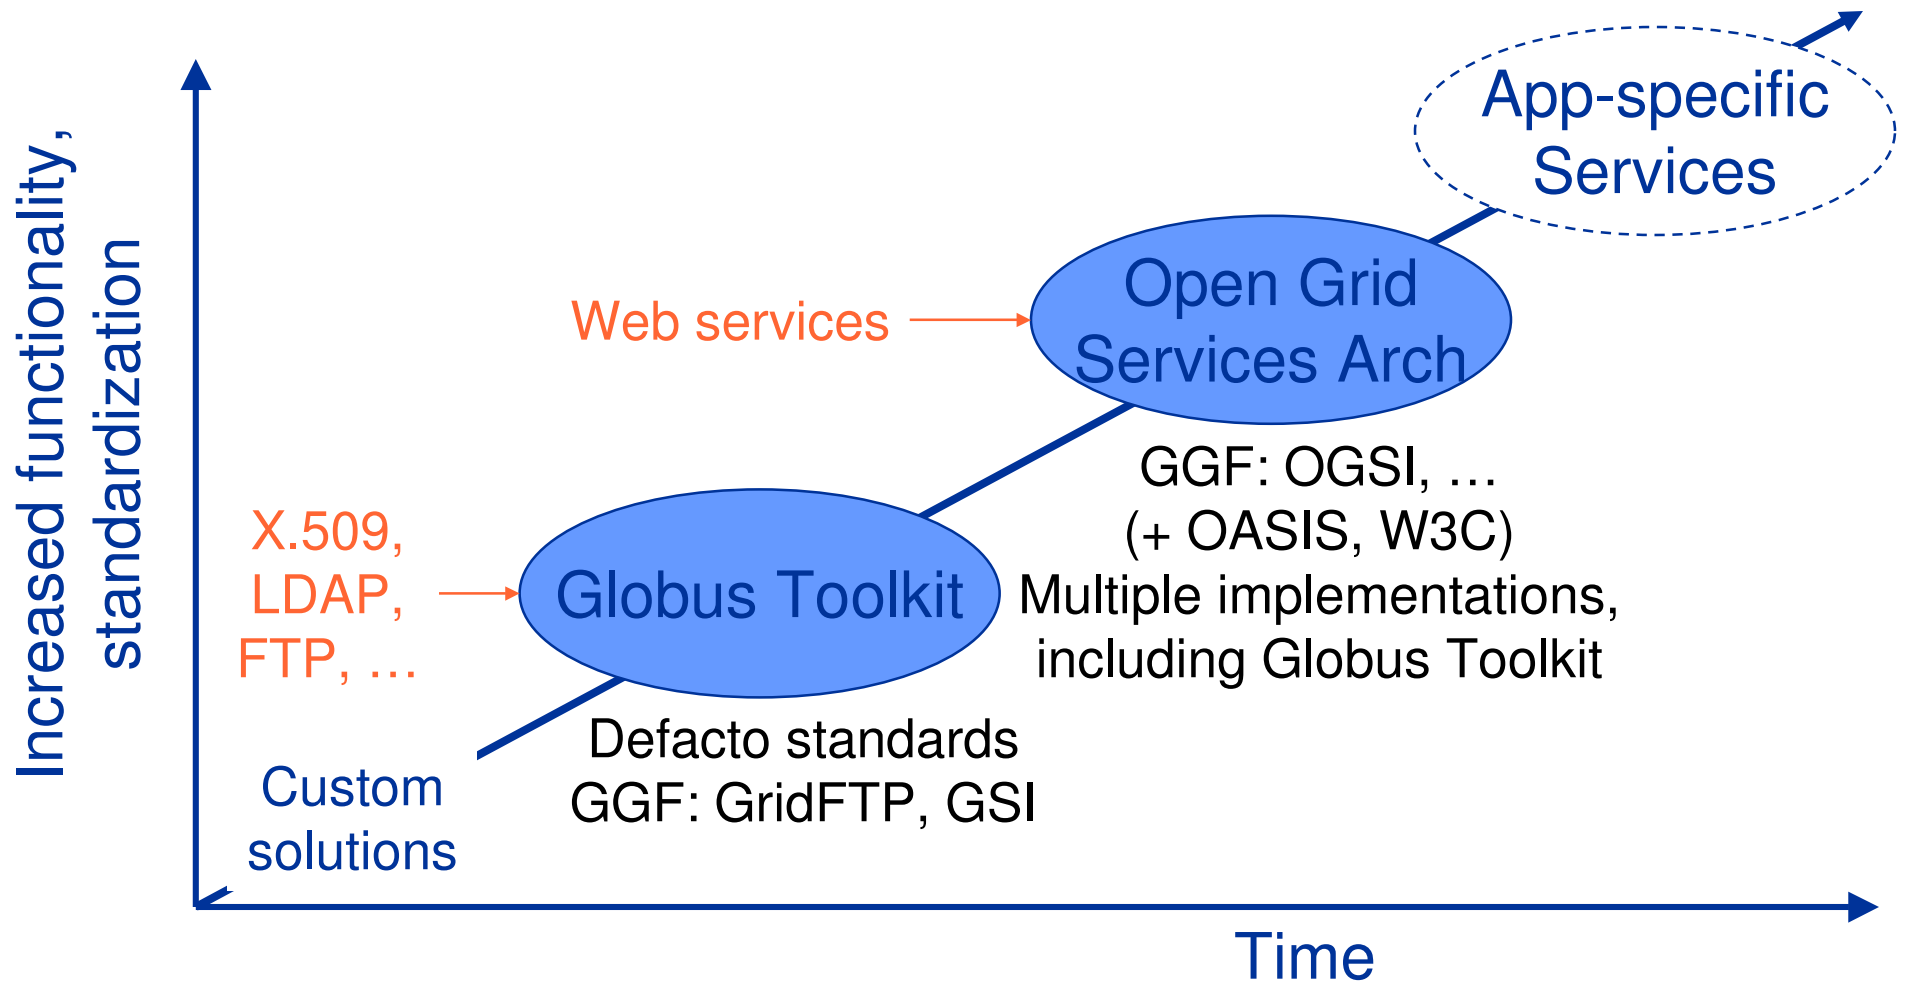
Realizing the Promise: Building an Open Grid



Why Open Standards Matter

- Ubiquitous adoption demands open, standard protocols
 - ◆ Standard protocols enable *interoperability*
 - ◆ Avoid product/vendor lock-in
 - ◆ Enables innovation/competition on end points
- Further aided by open, standard APIs
 - ◆ Standard APIs enable *portability*
 - ◆ Allow implementations to port to different vendor platforms
- Internet and Web as exemplars

Grids and Open Standards



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

More Specifically

- A standard substrate: the Grid service
 - ◆ OGSI = Open Grid Service Infrastructure
 - ◆ Web services interfaces and behaviors that address key distributed system issues
- ... supports standard service specifications
 - ◆ Resource mgt, dbms, workflow, security, ...
 - ◆ Target of current & planned GGF efforts
 - ◆ OGSA wg defines "OGSA compliance"
- ... and arbitrary application-specific services based on these & other definitions

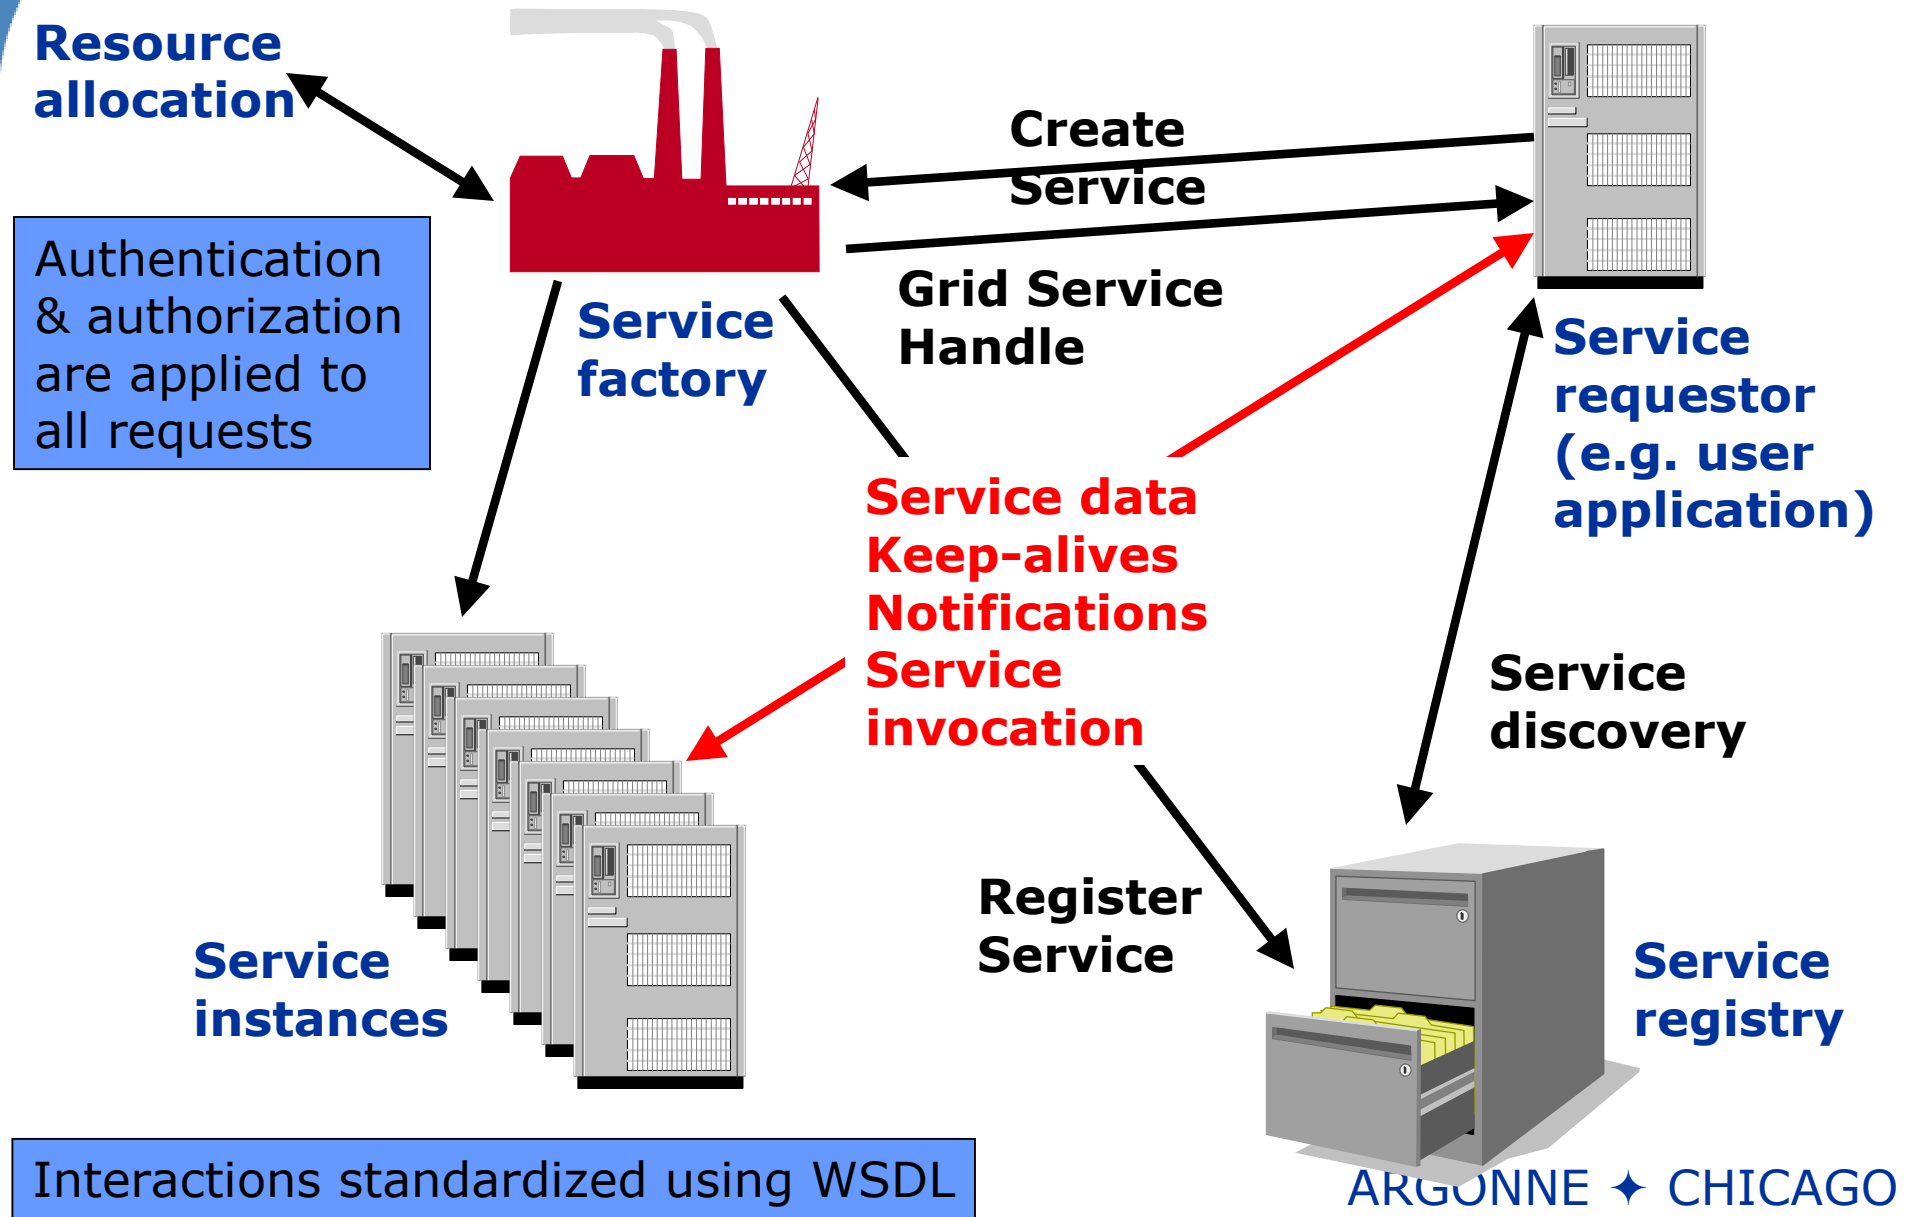
What Is OGSi?

- Useful, general purpose plumbing to make it easier to build Web services relevant to Grids
 - ◆ OGSi came about because we started trying to define Globus Toolkit® functionality using WSDL, and found there were common, base behaviors that we wanted to define once and reuse in all of our services
- There is nothing Grid specific about OGSi
- Perhaps it should have been better named:
WS-UsefulInterfacesToBuildAnInterestingClassOfWebServices

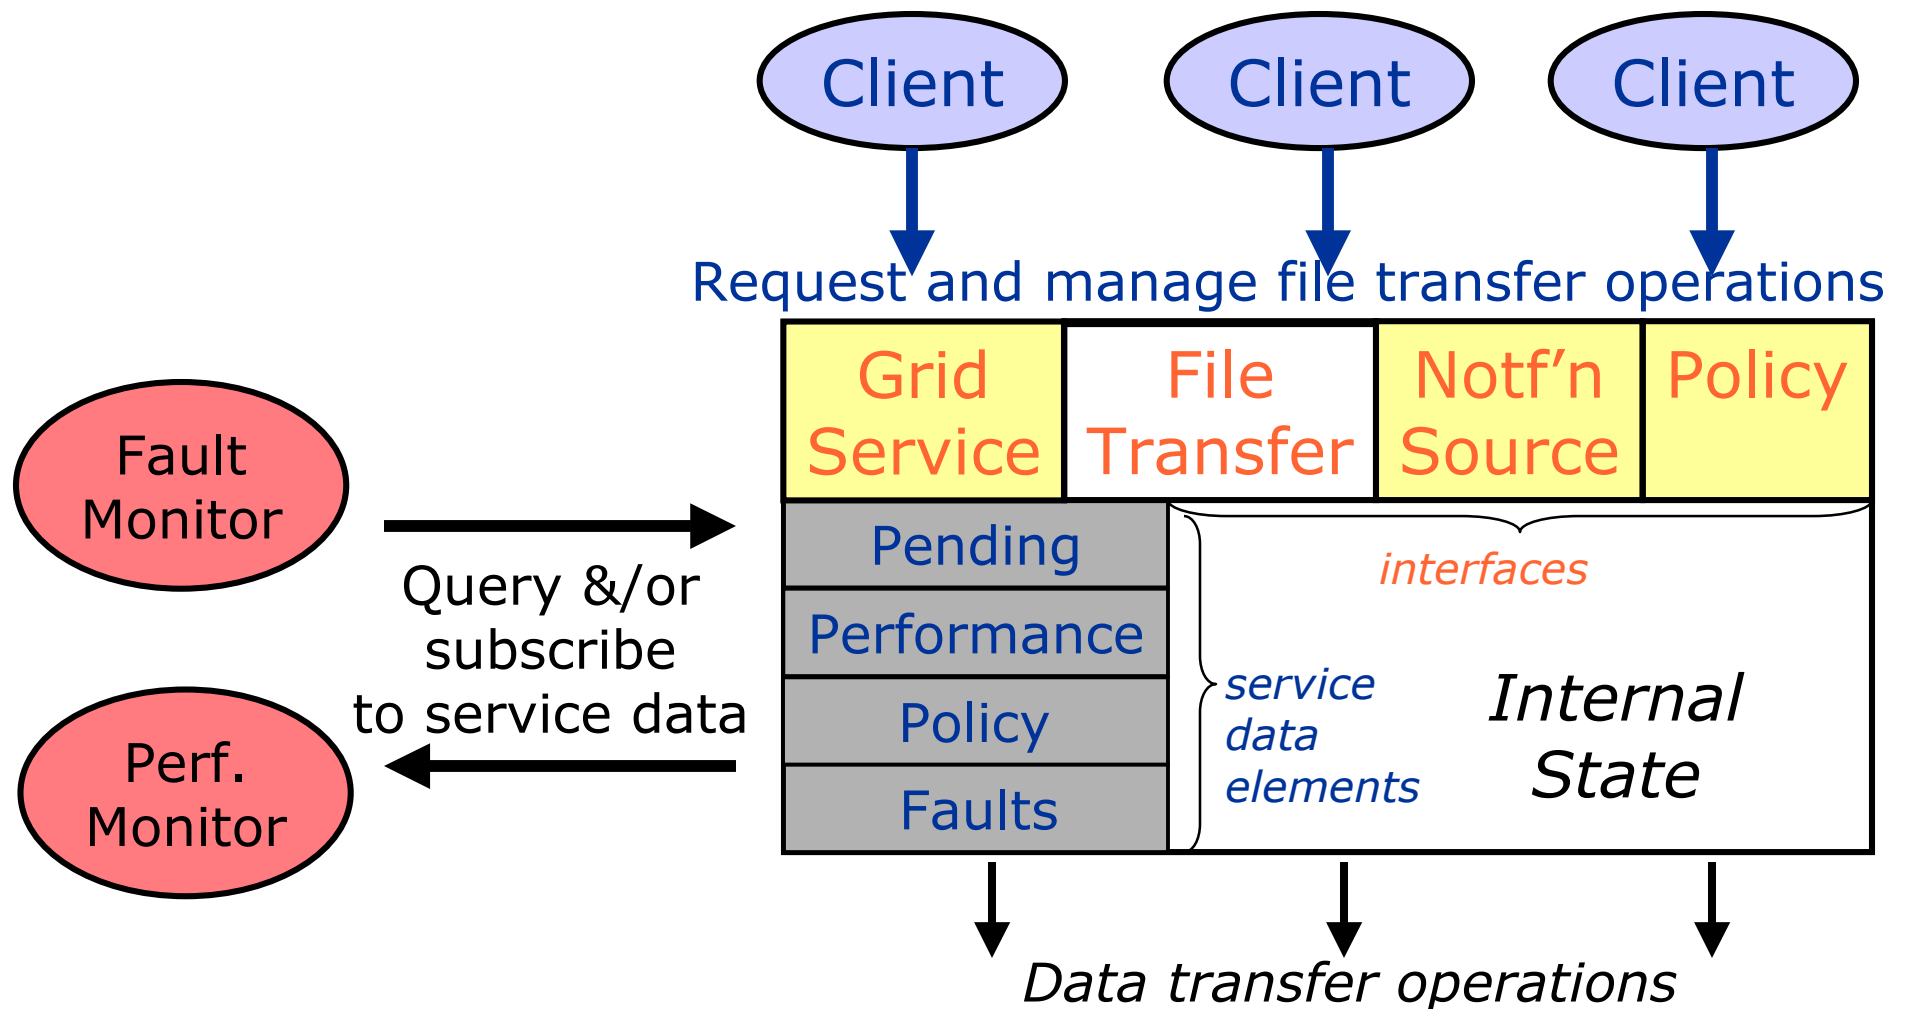
“Web Services”

- For OGSI, Web Services = WSDL, Web Services Description Language
 - ◆ OGSI is defined in terms of WSDL portTypes, messages, and XML Schema types
 - ◆ Largely silent on WSDL binding and service
- SOAP is important as a standard, inter-operable binding under WSDL, but OGSI is silent on this
 - ◆ Ditto for WS-Security, etc.
- UDDI registry can be used with OGSI
 - ◆ OGSI also defines primitives for building custom or domain-specific registries

Open Grid Services Infrastructure (OGSI)



Example: Reliable File Transfer Service



OGSA Standardization & Implementation

- OGSI defines core interfaces and behaviors for manageable services
 - ◆ Supported by strong open source technology & major commercial vendors
- Efforts are underway within GGF, OASIS, and other bodies to define standards for
 - ◆ Agreement negotiation
 - ◆ Common management model
 - ◆ Data access and integration
 - ◆ Security and policy
 - ◆ Etc., etc., etc.

Relevant Standards Organizations


- GGF: Grid services: OGSI/A, WS-Agreement
- W3C: Web services: WSDL, SOAP
- OASIS: Web services security, WSDM, SAML
- IETF: Internet protocols and security
- Project Liberty Alliance: Identity federation
- DMTF: Common Information Model (CIM)



Globus Toolkit Implements Standards As They Emerge

- GT 2.x
 - ◆ X.509 (Proxy) Certs, GridFTP, LDAP, GSS-API
- GT 3.0 (June 2003)
 - ◆ GT2 + WSDL, SOAP, OGSI, WS-Security, etc.
- GT 3.2 (1Q2004)
 - ◆ Maintenance release + new GridFTP code
- GT 3.x (3-4Q2004)
 - ◆ First implementation of many new OGSA standards: WS-Agreement, OGSA-DAI, ..., ...

Realizing the Promise: Building an Open Grid



Open Source

- Encourage adoption of standards by reducing barriers to entry
 - ◆ Overcome new technology Catch-22
- Enable broad Grid technology ecosystem
 - ◆ Key to international cooperation
 - ◆ Key to science-commerce partnership
- Jumpstart Grid industry and allow vendors to focus on value-add
 - ◆ E.g., IBM, Oracle, HP use GT; Platform Globus
 - ◆ Open source is industry friendly!

“Open Source” is More than Software and a License

- A community of contributors to code design, development, packaging, testing, documentation, training, support
 - ◆ United by common architectural perspective
- A community of users
 - ◆ May be major contributors via e.g. testing
- A governance structure
 - ◆ To determine how the software evolves
- Processes for coordinating all these activities
 - ◆ Packaging, testing, reporting, ...



The Globus Alliance & Toolkit

(Argonne, USC/ISI, Edinburgh, PDC)

- An international partnership dedicated to creating & disseminating high-quality open source Grid technology: the Globus Toolkit
 - ◆ Design, engineering, support, governance
- Academic Affiliates make major contributions
 - ◆ EU: CERN, Imperial, MPI, Poznan
 - ◆ AP: AIST, TIT, Monash
 - ◆ US: NCSA, SDSC, TACC, UCSB, UW, etc.
- Significant industrial contributions
- 1000s of users worldwide, many contribute

Other Relevant Efforts

- NSF Middleware Infrastructure initiative
 - ◆ Funds system integrators for middleware
- U.K. Open Middleware Initiative
 - ◆ Archive, integration, testing, documentation
- EU DataGrid, GridLab, EGEE, GriPhyN, Earth Science Grid, NEESgrid, etc., etc.
 - ◆ Produce/integrate software
- U.S. Virtual Data Toolkit
 - ◆ GT + Condor in physics-friendly package



Globus Toolkit Contributors Include

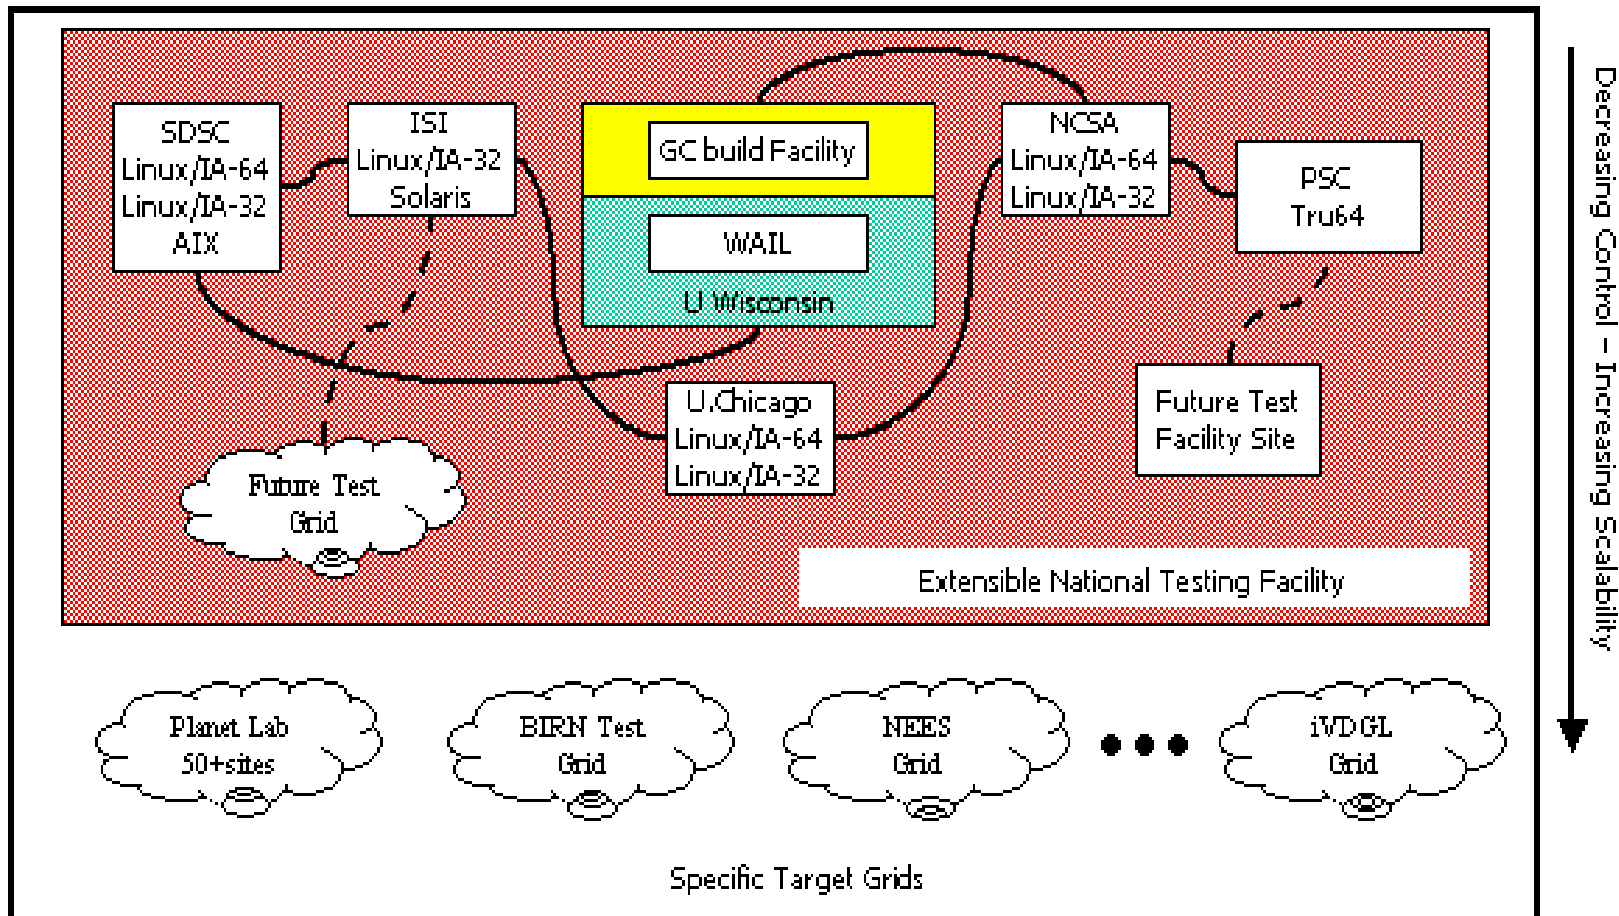
- Grid Packaging Technology (GPT) NCSA
- Persistent GRAM Jobmanager Condor
- GSI/Kerberos interchangeability Sandia
- Documentation NASA, NCSA
- Ports IBM, HP, Sun, SDSC, ...
- MDS stress testing EU DataGrid
- Support IBM, Platform, UK eScience
- Testing and patches Many
- Interoperable tools Many
- Replica location service EU DataGrid
- Python hosting environment LBNL
- Data access & integration UK eScience
- Data mediation services SDSC
- Tooling, Xindice, JMS IBM
- Brokering framework Platform
- Management framework HP
- \$\$ DARPA, DOE, NSF, NASA, Microsoft, EU
26 ARGONNE ♦ CHICAGO



GT Processes Are Evolving Rapidly (Thanks to Partners & the GT Team)

- Before 2000
 - ◆ Email-based problem tracking, aka “req”
- 2000
 - ◆ Detailed documentation, release notes (Q1)
 - ◆ Legal framework for external contributions (Q1)
- 2001
 - ◆ Packaging; module & binary releases (Q4)
 - ◆ Substantial regression tests (Q4)
- 2002
 - ◆ Bugzilla problem reporting & tracking (Q2)
 - ◆ Processes for external contrib (Q2)
 - ◆ Distributed testing infrastructure (Q3)
- 2003-2004 (in progress)
 - ◆ Distributed support infrastructure: GT “support centers”
 - ◆ Standardized Grid testing framework(s)
 - ◆ GT “contrib” components
 - ◆ Grid Technology Repository

Distributed Test Facility



(NSF GRIDS Center activity: grids-center.org)

the globus project
www.globus.org

Welcome to the Grid Technology Repository

Friday, January 10 2003 @ 06:41 PM CST

[advanced search](#) [Contact](#)

About GTR

Welcome to GTR!
This is a site devoted to the collection of OGSI-compliant Gridservices. You may submit your own using the "Get Published" Link, or download, comment, and vote on the works of others. The top-rated services are available from the "Voter's Picks" block.

Categories

Home
URIs (0/0)
Documentation (1/0)
Code (2/0)

User Functions

Username:

Password:

Login

Don't have an account yet?
Sign up as a [New User](#)

20 Most Recently Posted [See Full List]

Contributor	Contribution Name	Description
bacon	A new gridservice	This is not really a gridservice.
Admin	new code	some new code
Admin	new doc	some new doc for the db

Top 10 Viewed Contributions

Contributor	Contribution Name	Description
Admin	new doc	some new doc for the db
bacon	A new gridservice	This is not really a gridservice.
Admin	new code	some new code

Top 10 Highest Rated

Contribution Name	Contributor	Description	Votes / Avg
A new gridservice	bacon	code	1 / 5.00
new doc	Admin	doc	2 / 4.00
new code	Admin	code	1 / 1.00

<http://gtr.globus.org>


The Grid Technology Repository

- Community repository
- Clearing house for service definitions, code, documentation
- Encourage collaboration & avoid redundant work

International advisory committee: Ian Foster (Chair), Malcolm Atkinson, John Brooke, Fabrizio Gagliardi, Dennis Gannon, Wolfgang Gentzsch, Andrew Grimshaw, Keith Jackson, Gregor von Laszewski, Satoshi Matsuoka, Jarek Nabrzyski, Bill St. Arnaud, Jay Unger

20+ contributions received

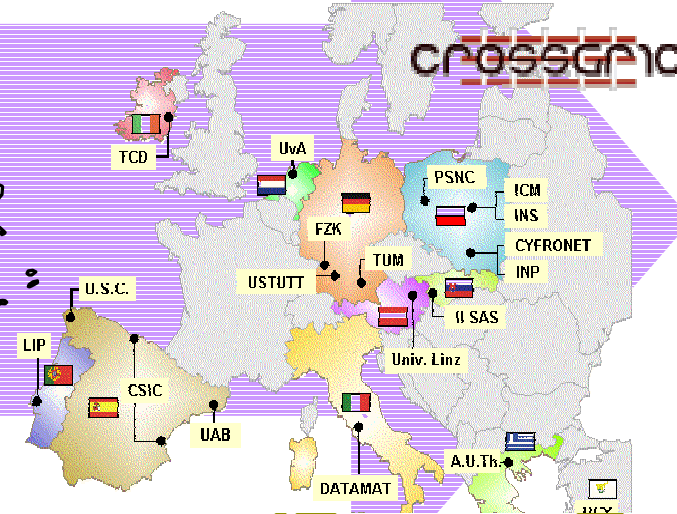
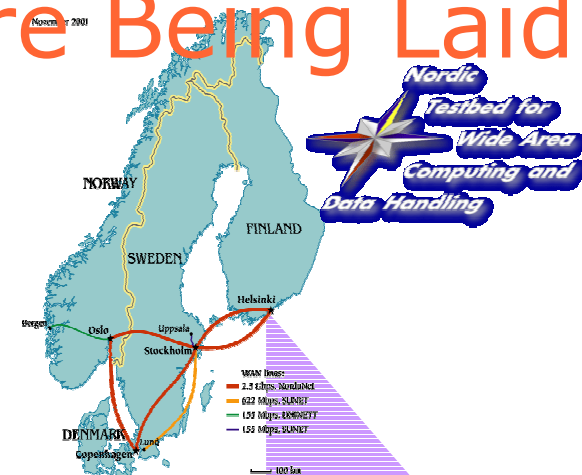

Realizing the Promise: Building an Open Grid



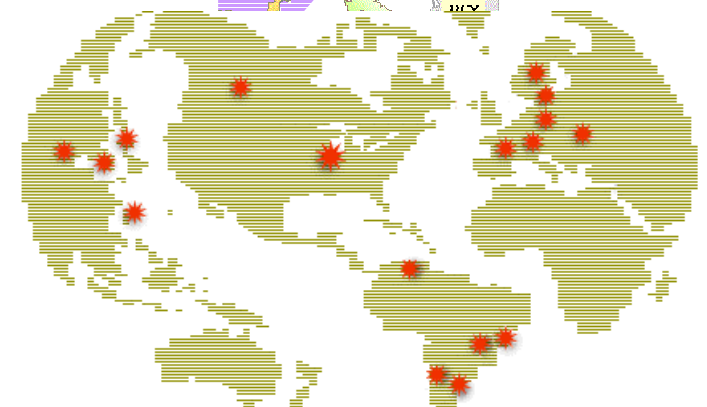
Grid Infrastructure

- Broadly deployed services in support of fundamental collaborative activities
 - ◆ Formation & operation of virtual organizations
 - ◆ Authentication, authorization, discovery, ...
- Services, software, and policies enabling on-demand access to critical resources
 - ◆ Computers, databases, networks, storage, software services,...
- Operational support for 24x7 availability
- Integration with campus and commercial infrastructures

The Foundations Are Being Laid

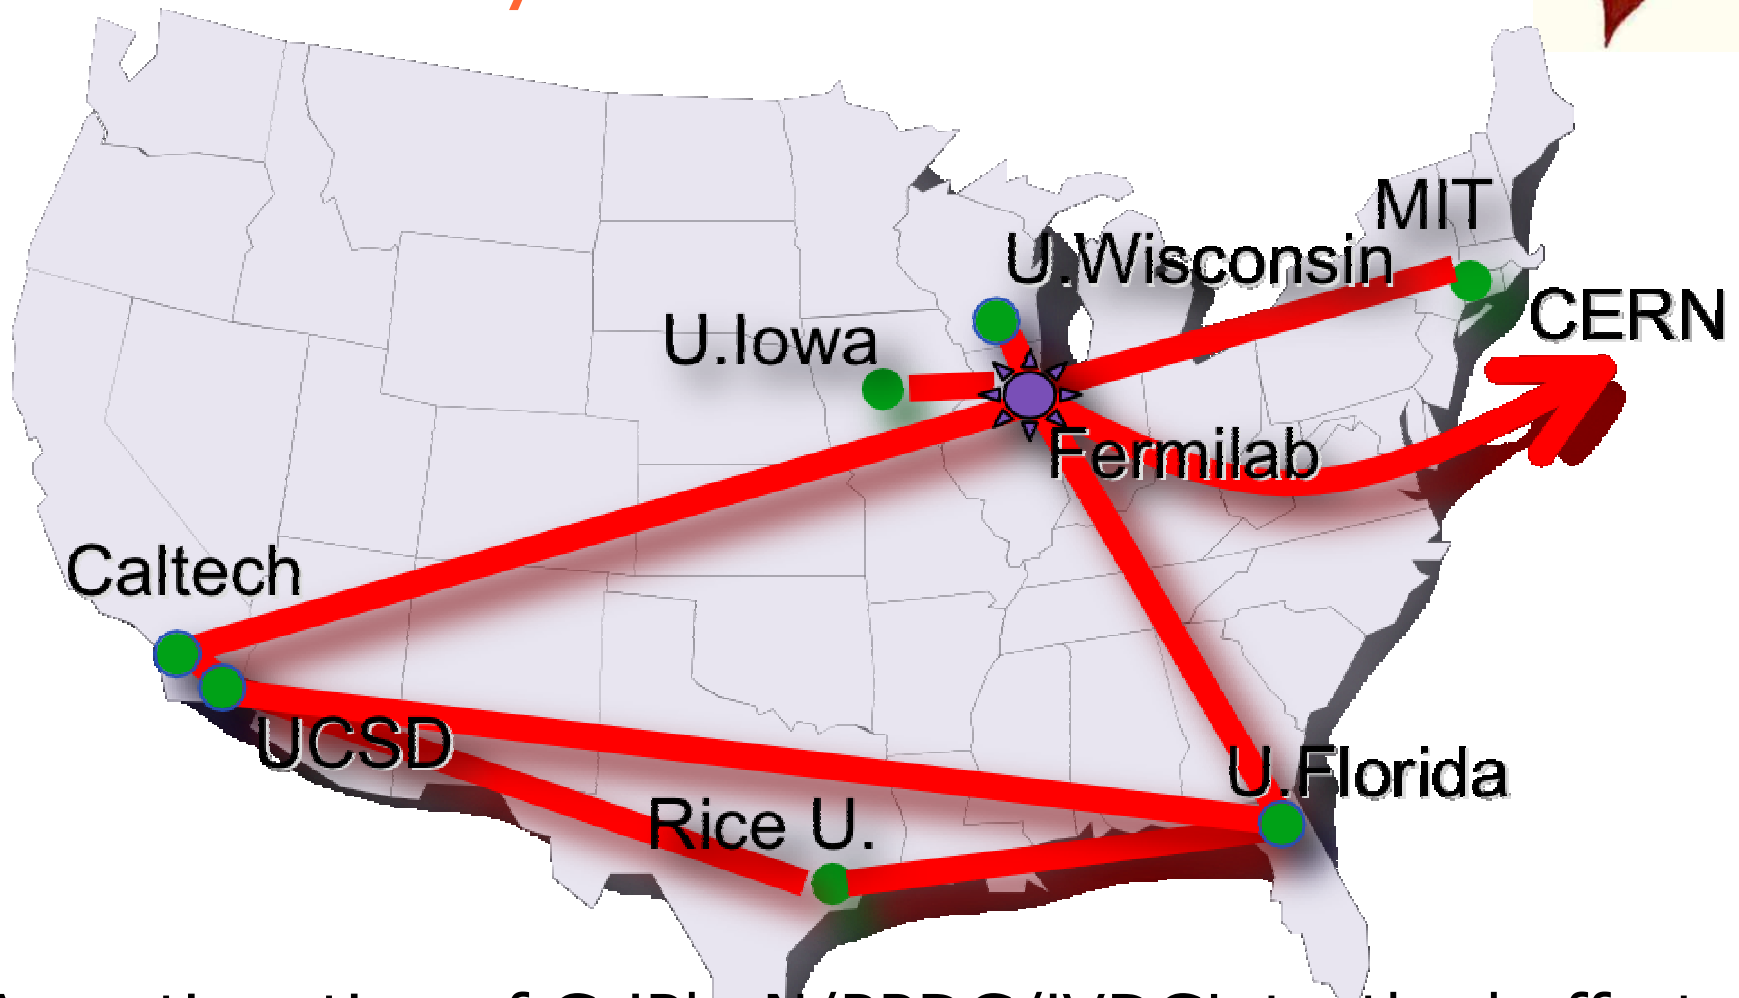


TERAGRID



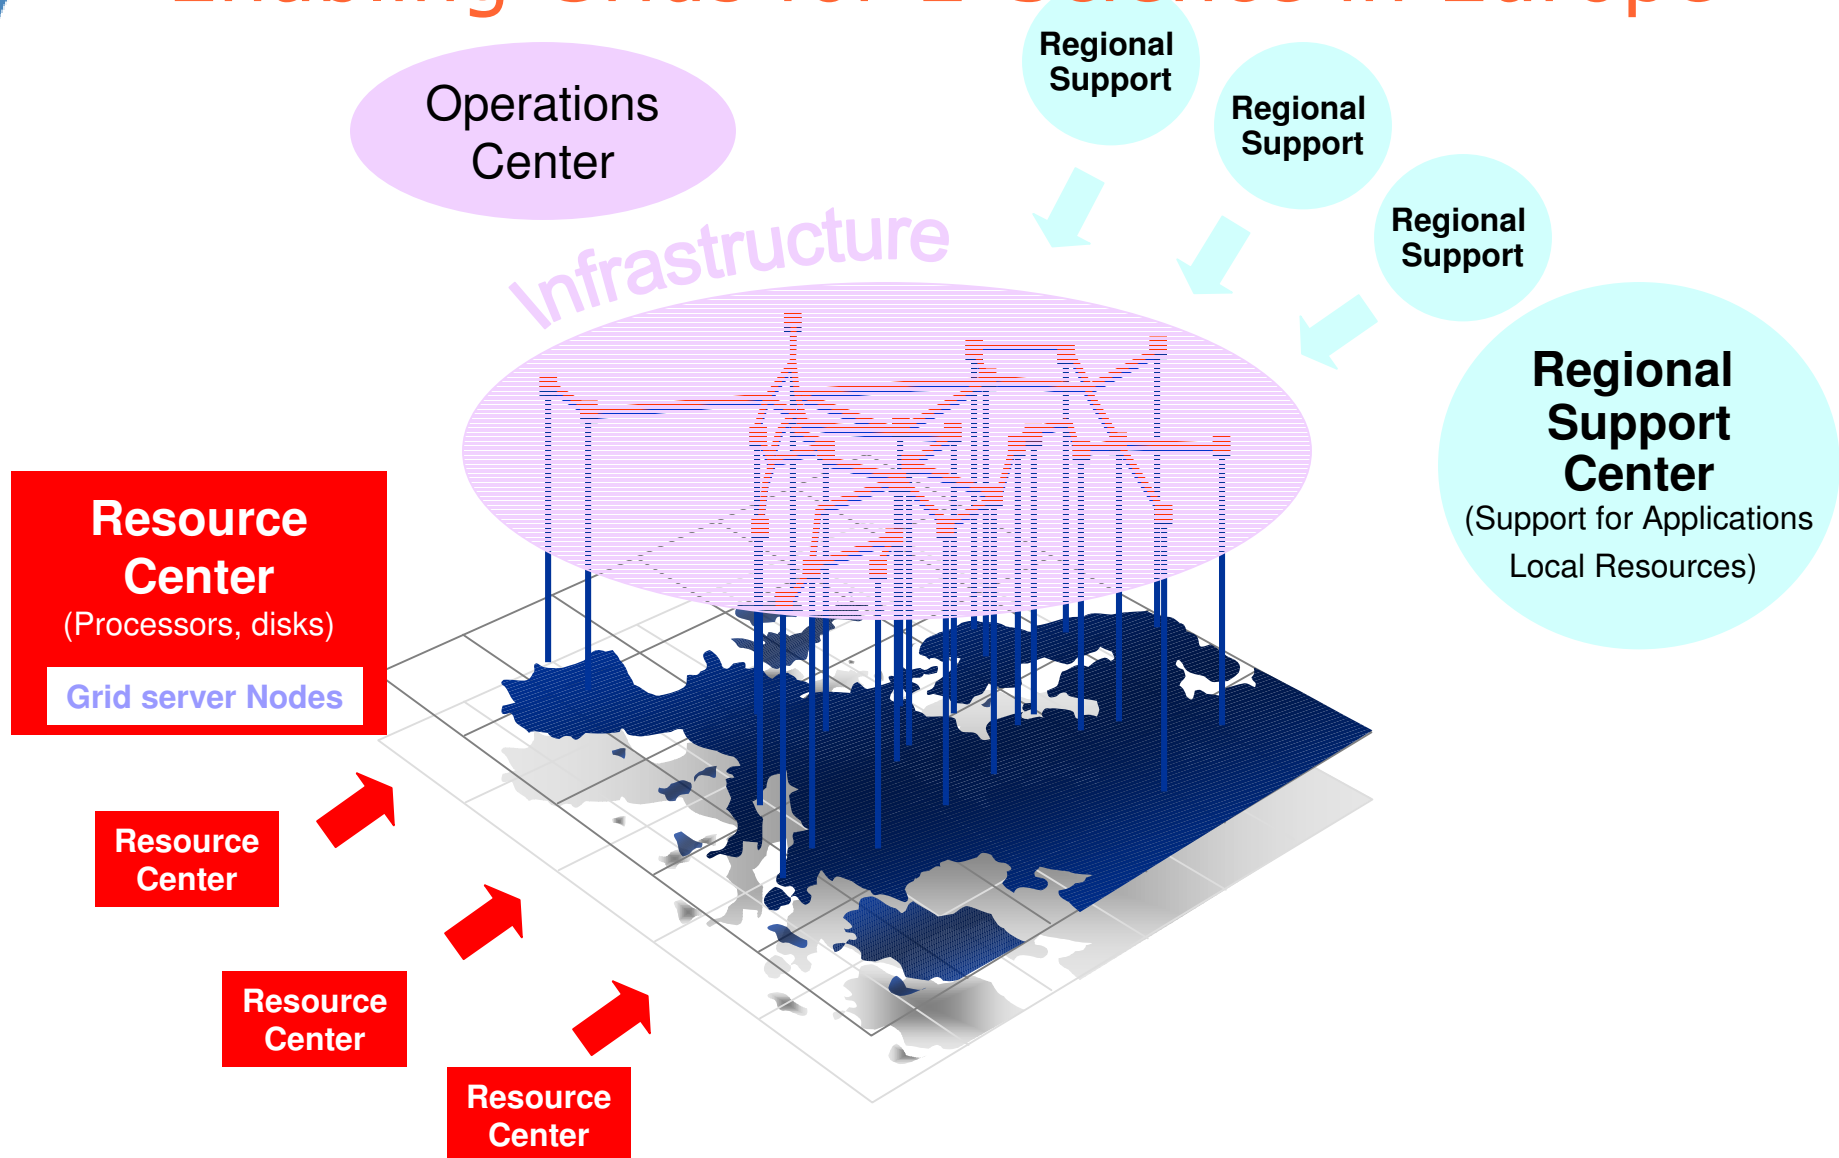


Current Focus of U.S. Physics Grid Efforts: Grid3




A continuation of GriPhyN/PPDG/iVDGL testbed efforts,
focused on establishing a functional federated Grid

EGEE: Enabling Grids for E-Science in Europe



Realizing the Promise: Building an Open Grid



What We Can Do Today

- A core set of Grid capabilities are available and distributed in good quality form, e.g.
 - ◆ GT: security, discovery, access, data movement
 - ◆ Condor, EDG: scheduling, workflow management
 - ◆ Virtual Data Toolkit, NMI, etc.
- Deployed at moderate scales
 - ◆ NEESgrid, EU DataGrid, Grid3, LCG-1, ...
- Usable with some hand holding, e.g.
 - ◆ CMS event production: 5 sites, 2 months
 - ◆ NEESgrid earthquake engineering experiment



Back Forward Stop Refresh Home Print Mail Add

Address <http://grid.tacc.utexas.edu/>

TACC ATT Grid TDR WAMI Google UT Portal



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](#)

	System/ Processors	Peak GFLOPs	Memory GBytes	Work Disk GBytes	Name	Grid SW	Network	Status	Load	Jobs
CS	Linux PC	1.8	1	92	alta	Q		↑		
CS	Linux PC	1.5	1	52	solitude	Q		↑		
TACC	Cray SV1 / 16	19	16	465	aurora	Q		↑		78
TACC	Linux Cluster / 2	1	6	13	brason	Q		↑		
TACC	Linux PC	2	1	10	cool	Q	Q	↑		
TACC	IBM Regatta-HPC / 64	303	128	532	longhorn	Q		↑		4R-4Q
TACC	LSF Multi-Clustered 22	37	14	173	lal	Q		↑		6R-2Q-3Q
TACC	Linux Cluster / 4	2	1	13	padre	Q		↑		
TACC	Cray/Dell Cluster / 4	19	8	8	q	Q		↑		
TACC	Linux PC	2	1	10	sanantonio	Q	Q	↑		
TACC	IBM IA-64 Cluster / 40	128	80	140	santarita	Q	Q	↑		
TACC	Sun Workstation	2	1	2	tahoka	Q		↑		
TACC	IBM IA-32 Cluster / 64	64	32	20	tejas	Q	Q	↑		6R-4Q-3Q
TACC	Alpha Cluster / 16	10	8	71	zaphod	Q		↑		
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.



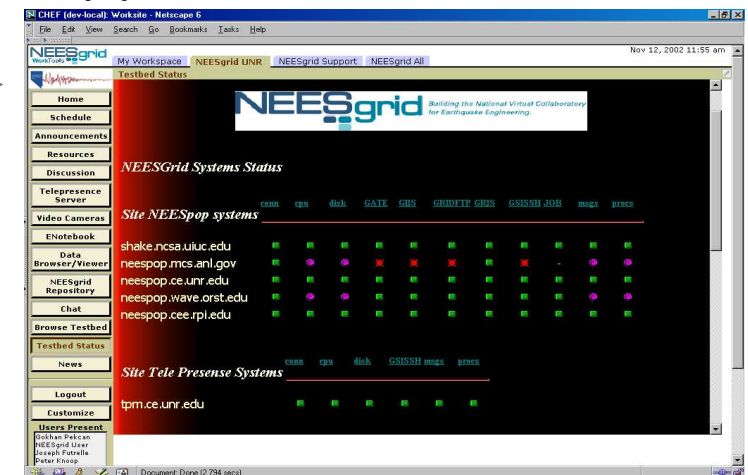
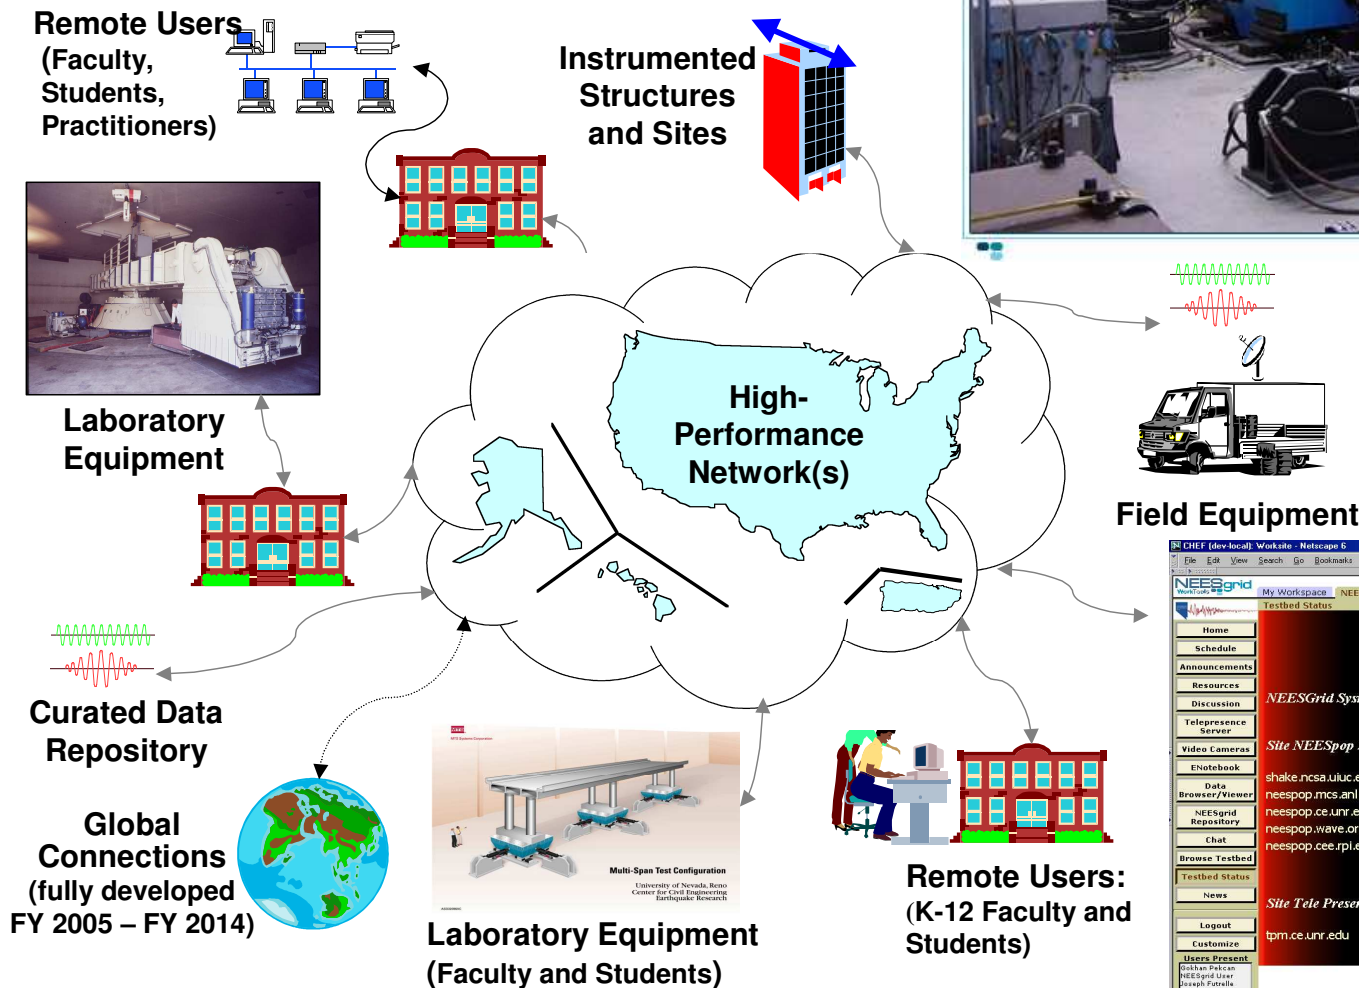
the globus alliance
www.globus.org

NEESgrid Earthquake Engineering Collaboratory



U.Nevada Reno

www.neesgrid.org



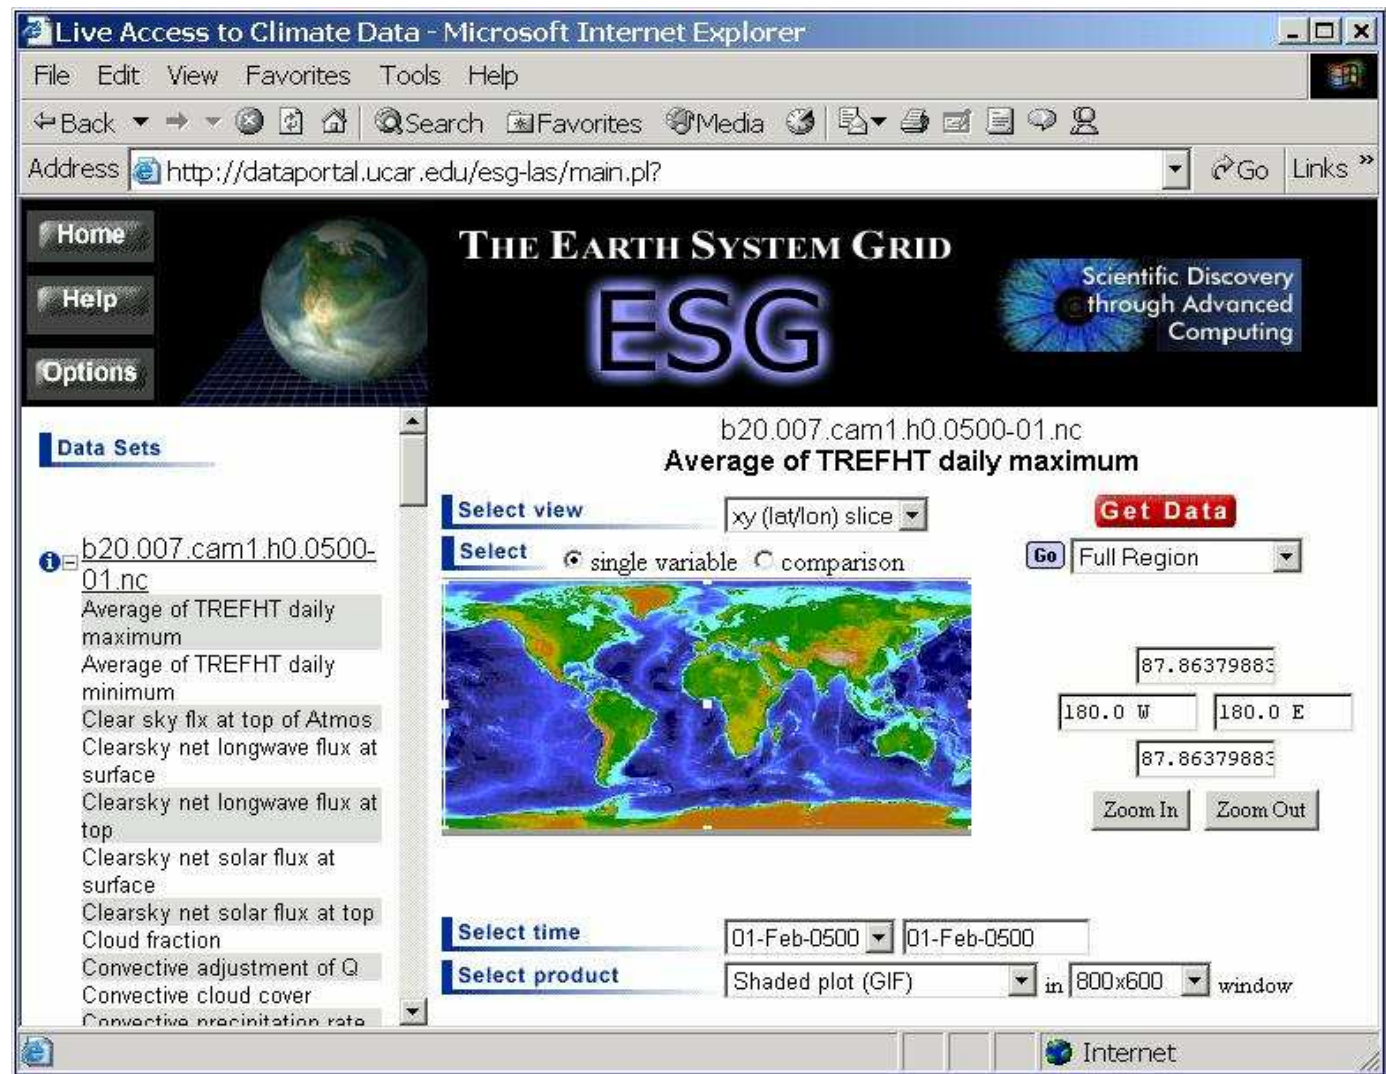


the globus alliance

www.globus.org

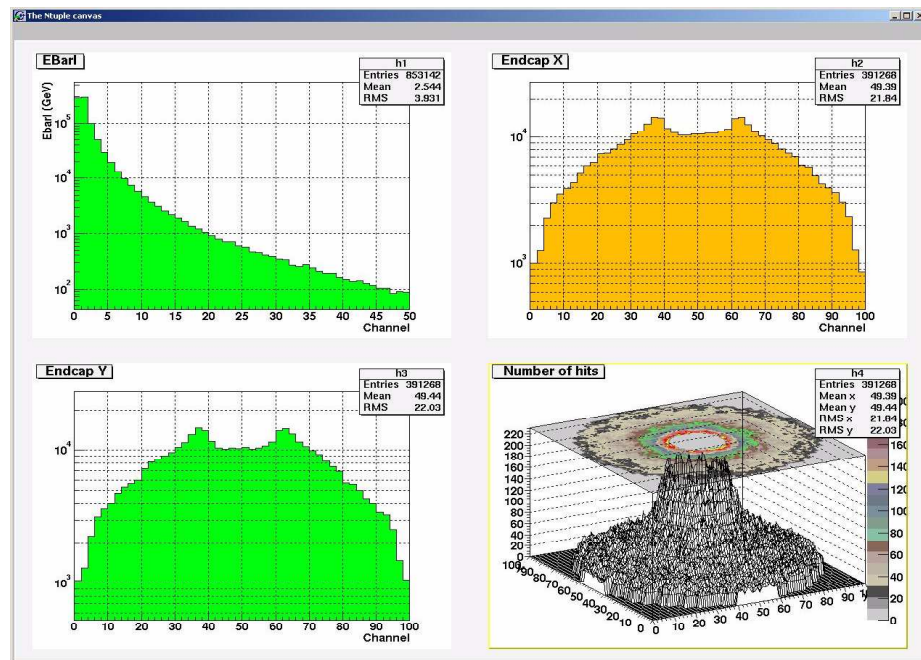

Earth System Grid (ESG)

Goal: address technical obstacles to the sharing & analysis of high-volume data from advanced earth system models



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

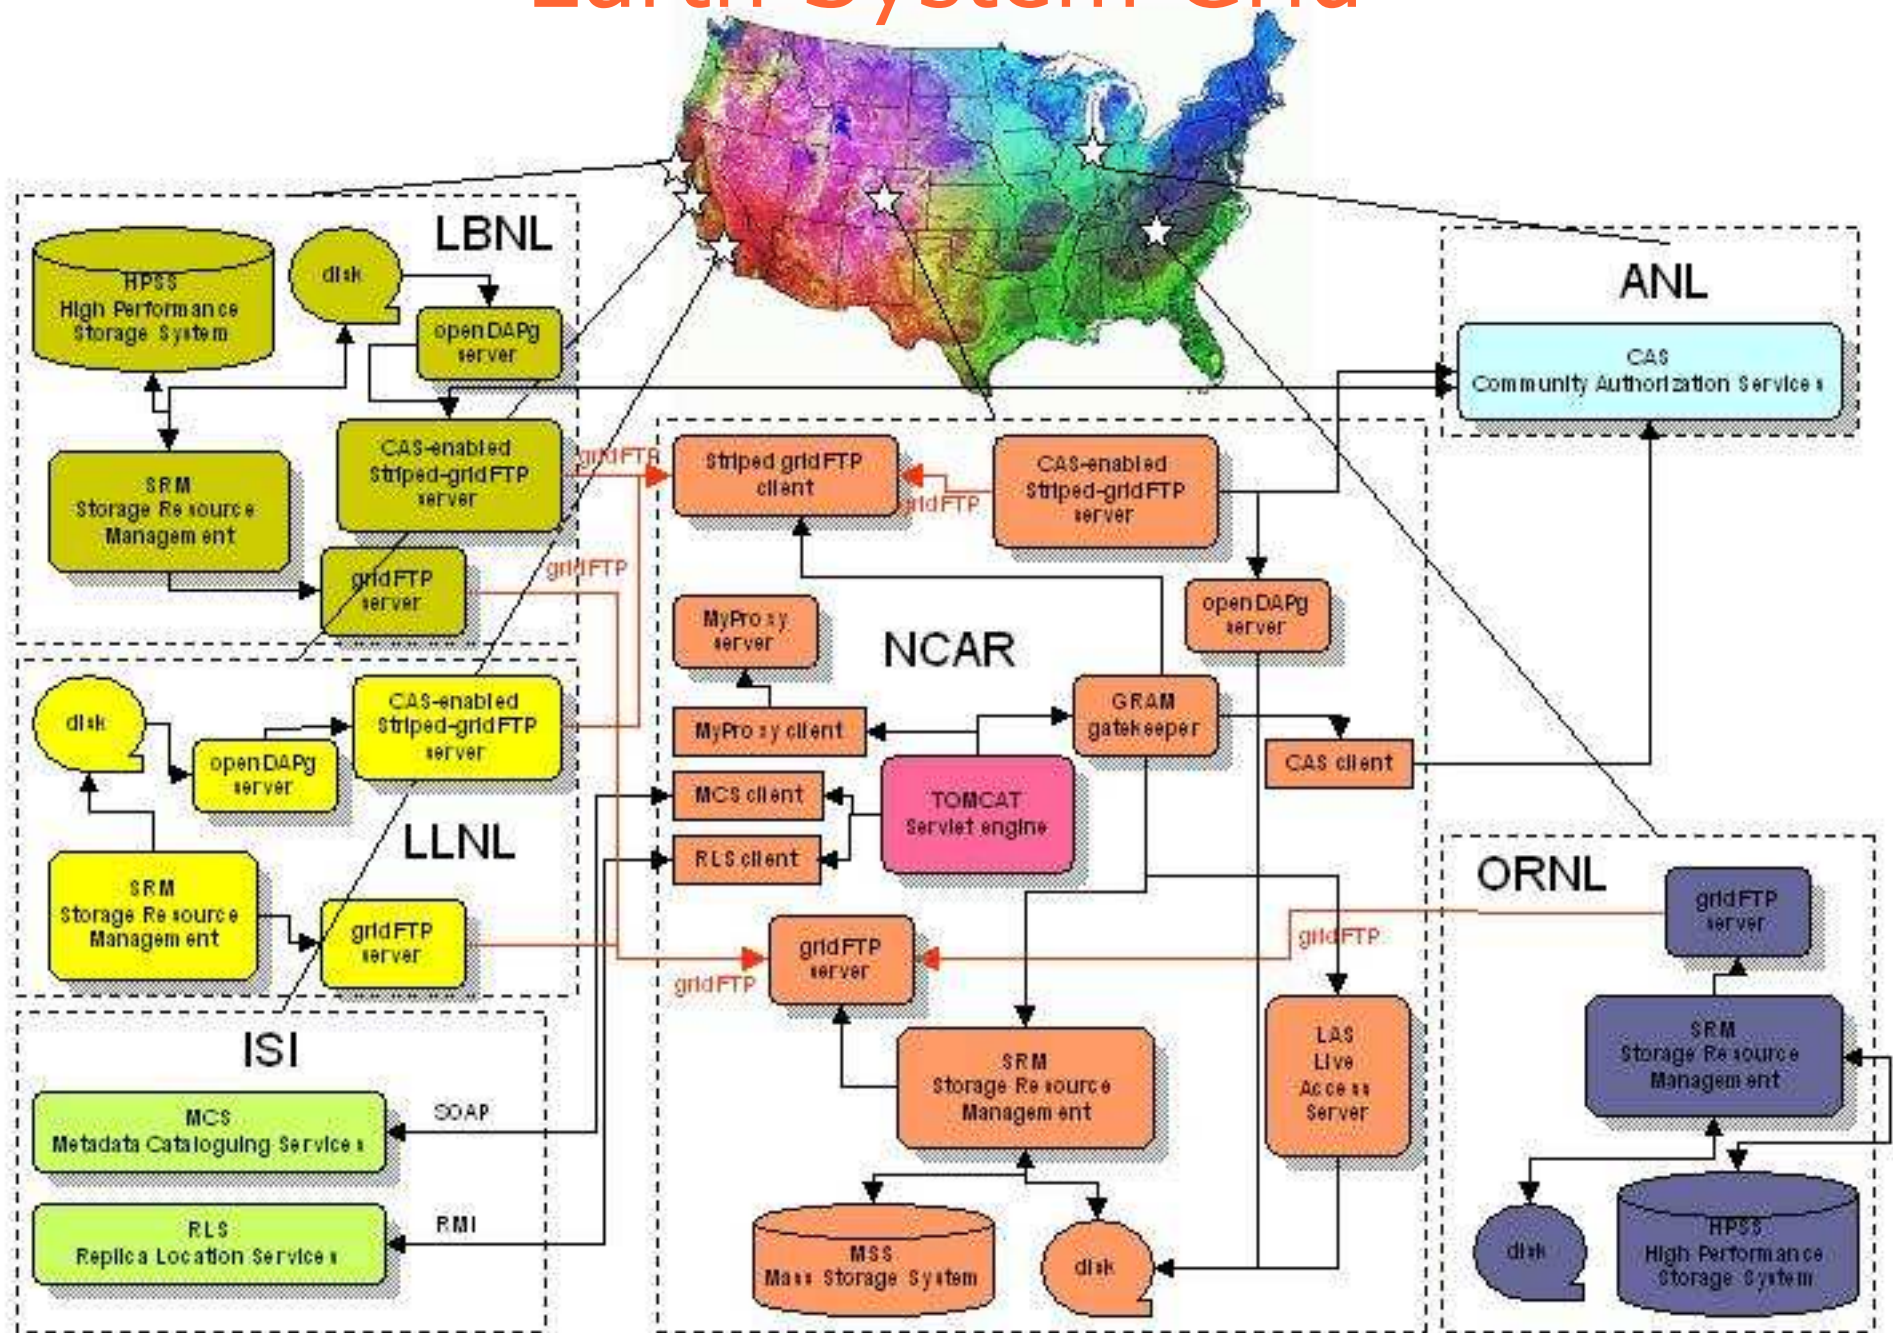


- EU DataGrid and LCG-1 operating at similar scales

Challenges

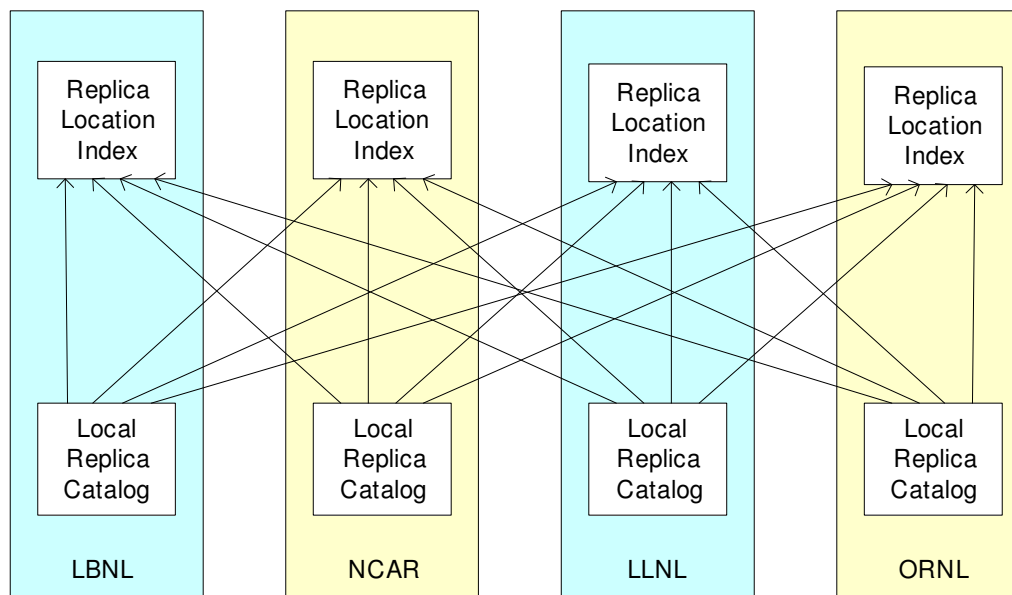
- Integration with site operational procedures
 - ◆ Many difficult issues
- Scalability in multiple dimensions
 - ◆ Number of sites, resources, users, tasks
- Higher-level services in multiple areas
 - ◆ Virtual data, policy, collaboration
- Integration with end-user science tools
 - ◆ Science desktops
- Coordination of international contributions
- Integration with commercial technologies

Earth System Grid



Scalable, Resilient Services: EDG/Globus Replica Location

- Soft state maintenance of replicated index
- E.g., Earth System Grid replica registry



Strategies for Success

- Establish real international collaboration
 - ◆ Partition the space of what to do: it's large
 - ◆ Be each other's partners, not customers
- Tackle process as well as software
 - ◆ Standardize on packaging, testing, support
 - ◆ Deployment and operations issues
- Build sustained, critical mass teams
 - ◆ Software is hard; requires time & expertise
- Build and operate large-scale Grids with real application groups to drive all of this
 - ◆ Look beyond data grids and physics

Strategies for Failure

- Select software for political reasons
 - ◆ E.g., “We must have a European solution”
- Underestimate cost & difficulty of software
 - ◆ “It’s easy to hire programmers & students”
- Start from scratch if “A does not do Z”
 - ◆ You don’t need U, V, W, X, Y. Or do you?
- Embrace diversity and non-interoperability
 - ◆ “We can add an interoperability layer”
- Develop discipline-specific solutions

Summary

- We have made a good start on creating an open Grid technology in three key areas
 - ◆ Standards, software, infrastructure
 - ◆ Real application success stories
- We must now scale this technology base
 - ◆ Address deployment & operational issues
 - ◆ Functionality and resilience at large scales
- International collaboration is key to success: we've scored a C so far, let's aim for an A

For More Information

- The Globus Alliance®
 - ◆ www.globus.org
- Global Grid Forum
 - ◆ www.ggf.org
- Background information
 - ◆ www.mcs.anl.gov/~foster
- **GlobusWORLD 2004**
 - ◆ www.globusworld.org
 - ◆ **Jan 20–23, San Francisco**

