U.S. Grid Projects: Grid3 and Open Science Grid

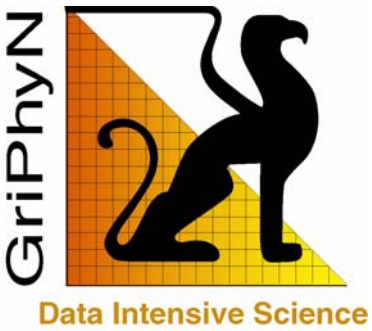

Paul Avery
University of Florida
avery@phys.ufl.edu

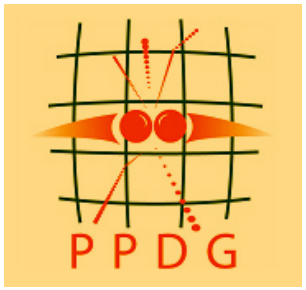

International ICFA Workshop on HEP, Networking & Digital Divide Issues for Global e-Science
Daegu, Korea
May 23, 2005

U.S. "Trillium" Grid Partnership

- Trillium = PPDG + GriPhyN + iVDGL
  - Particle Physics Data Grid: $12M (DOE) (1999 – 2006)
  - GriPhyN: $12M (NSF) (2000 – 2005)
  - iVDGL: $14M (NSF) (2001 – 2006)
- Basic composition (~150 people)
  - PPDG: 4 universities, 6 labs
  - GriPhyN: 12 universities, SDSC, 3 labs
  - iVDGL: 18 universities, SDSC, 4 labs, foreign partners
  - Expts: BaBar, D0, STAR, Jlab, CMS, ATLAS, LIGO, SDSS/NVO
- Coordinated internally to meet broad goals
  - GriPhyN: CS research, Virtual Data Toolkit (VDT) development
  - iVDGL: Grid laboratory deployment using VDT, applications
  - PPDG: "End to end" Grid services, monitoring, analysis
  - Common use of VDT for underlying Grid middleware
  - Unified entity when collaborating internationally

Goal: Peta-scale Data Grids for Global Science

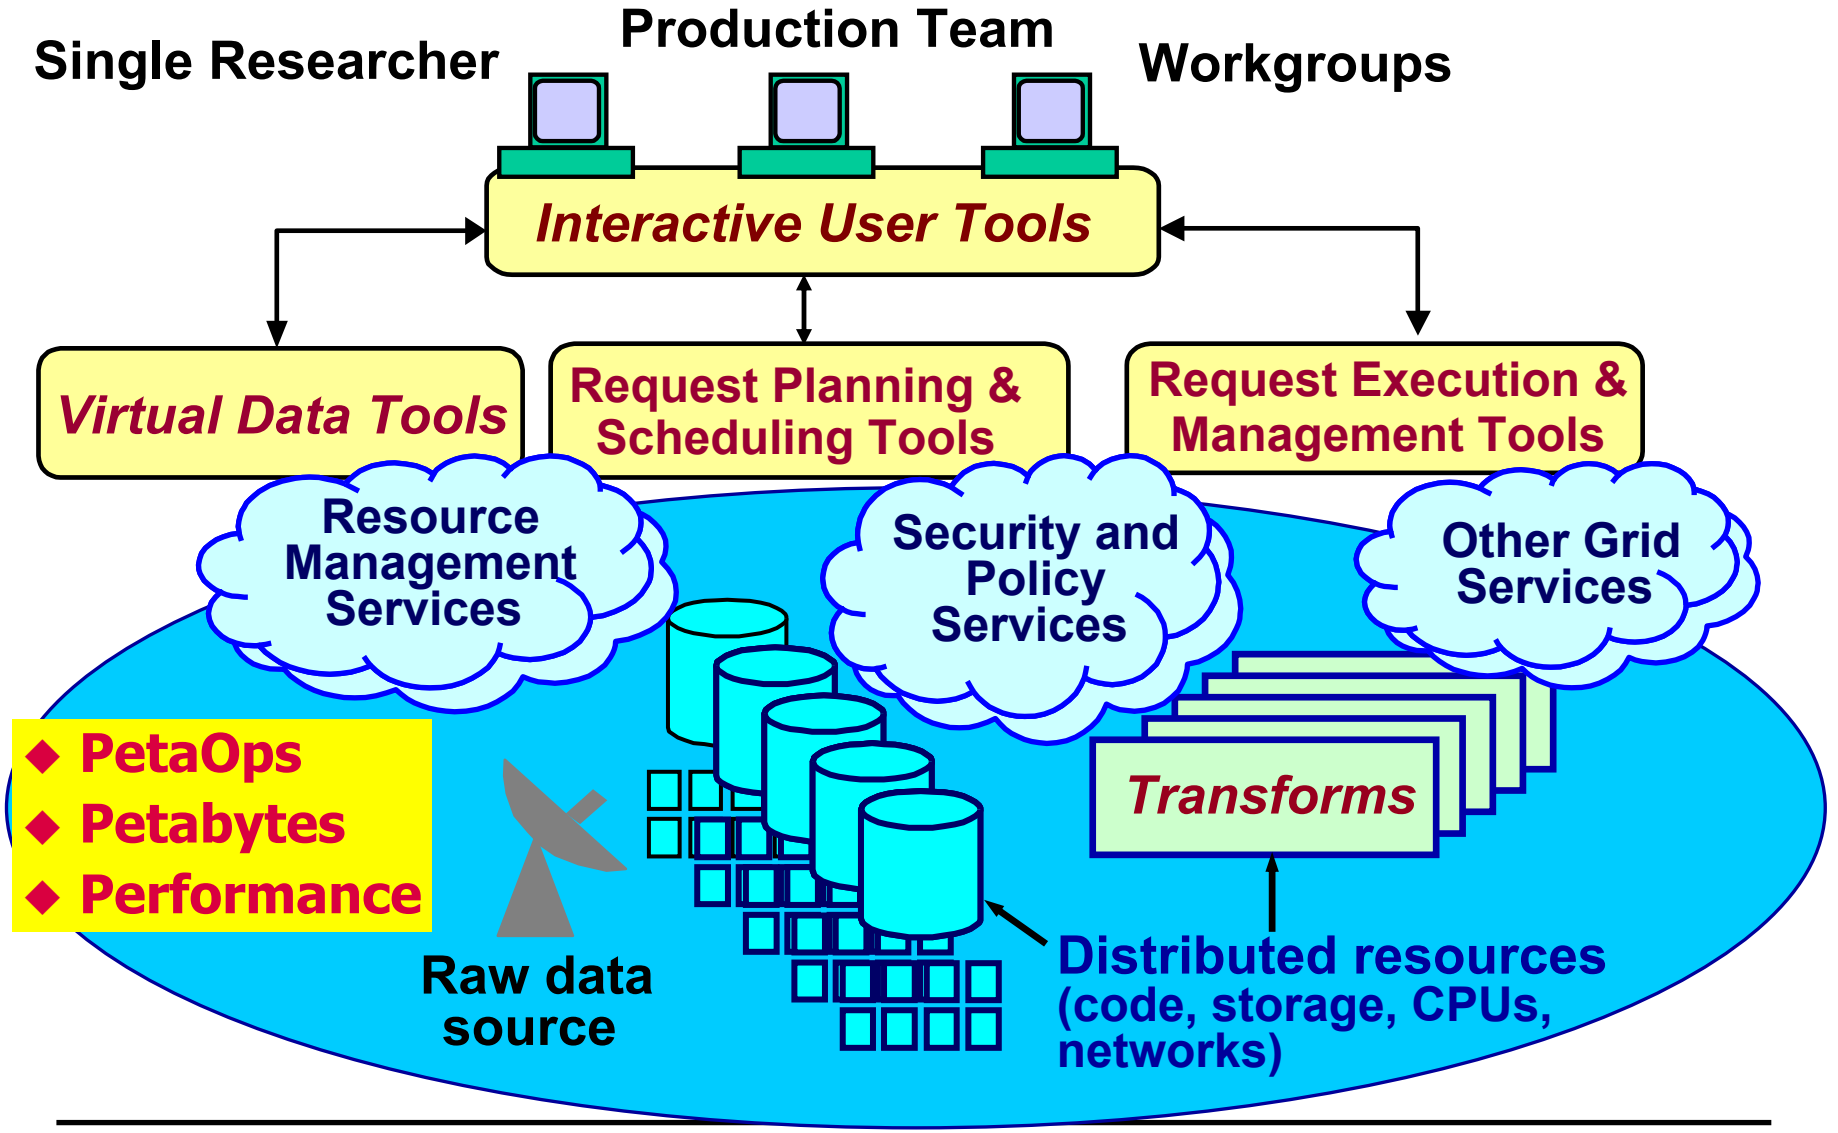

Grid Middleware: Virtual Data Toolkit

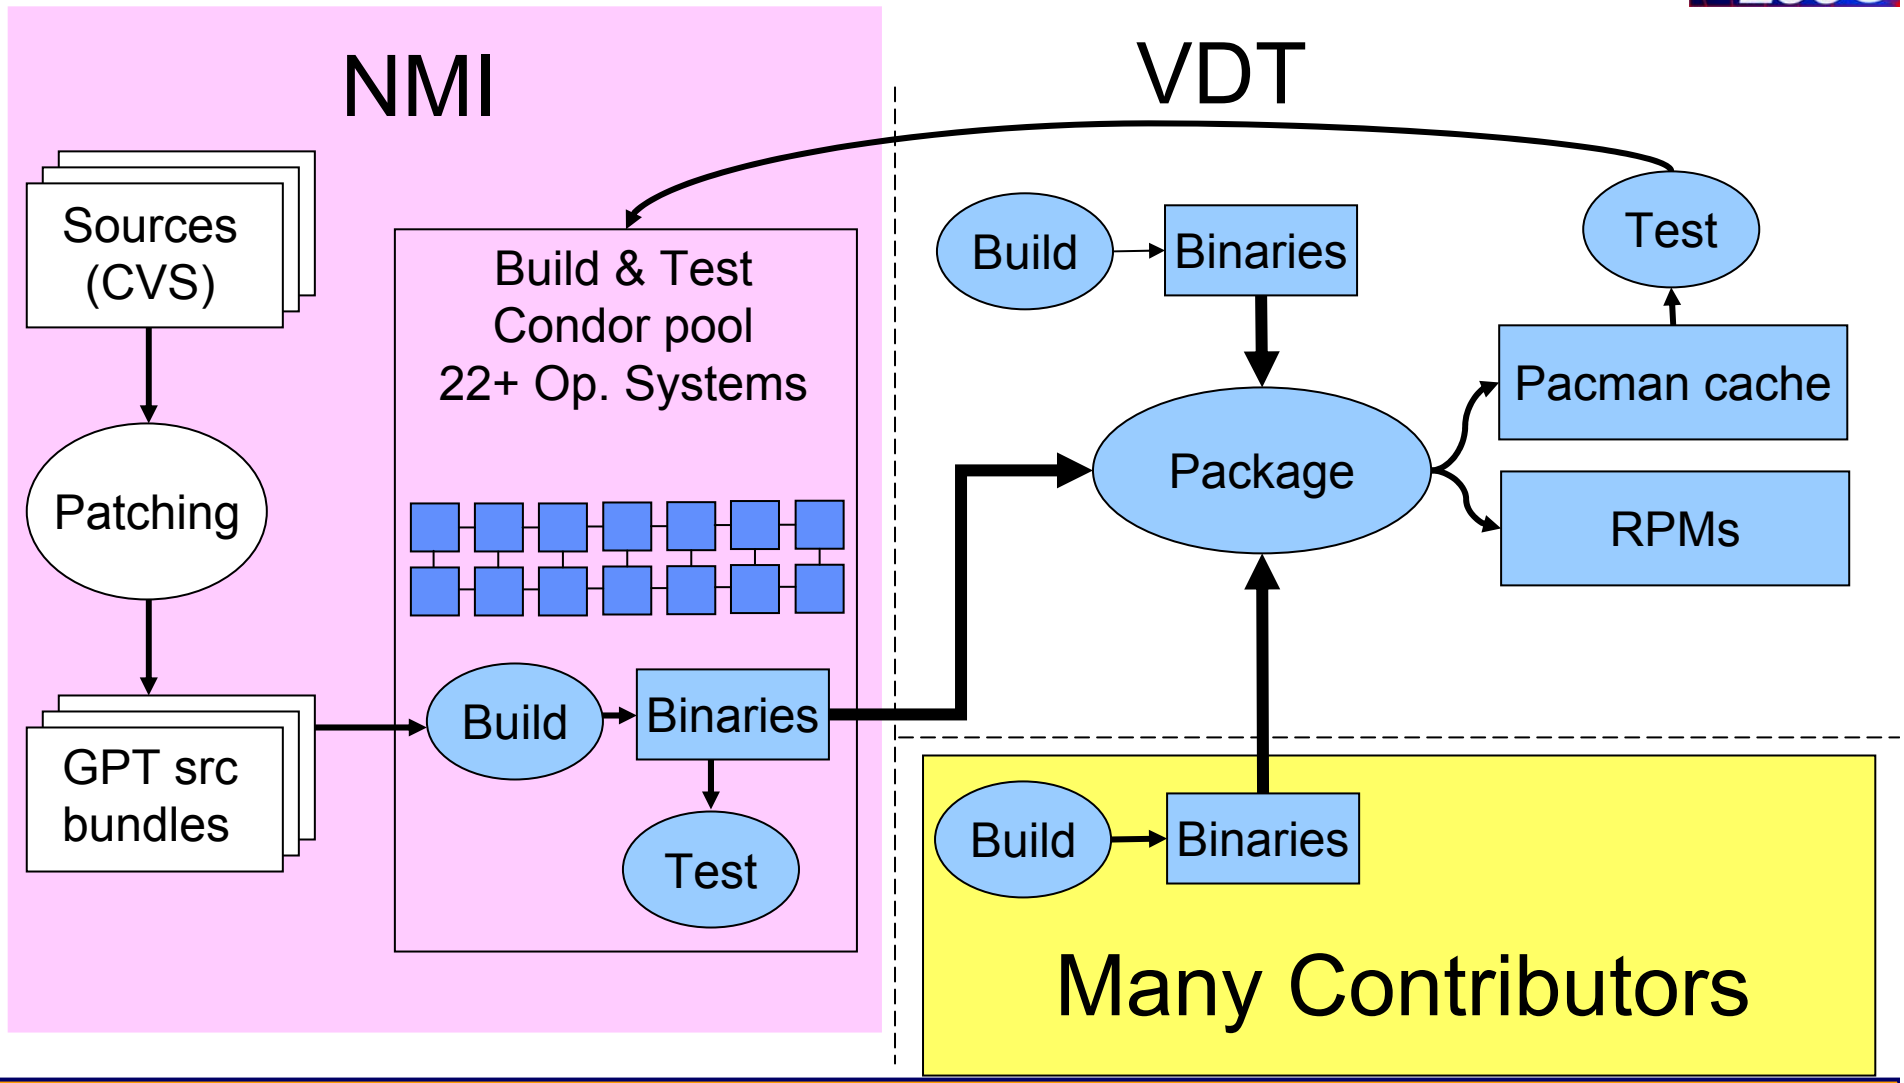

A unique laboratory for testing, supporting, deploying, packaging, upgrading, & troubleshooting complex sets of software!

VDT Growth Over 3 Years

www.griphyn.org/vdt/

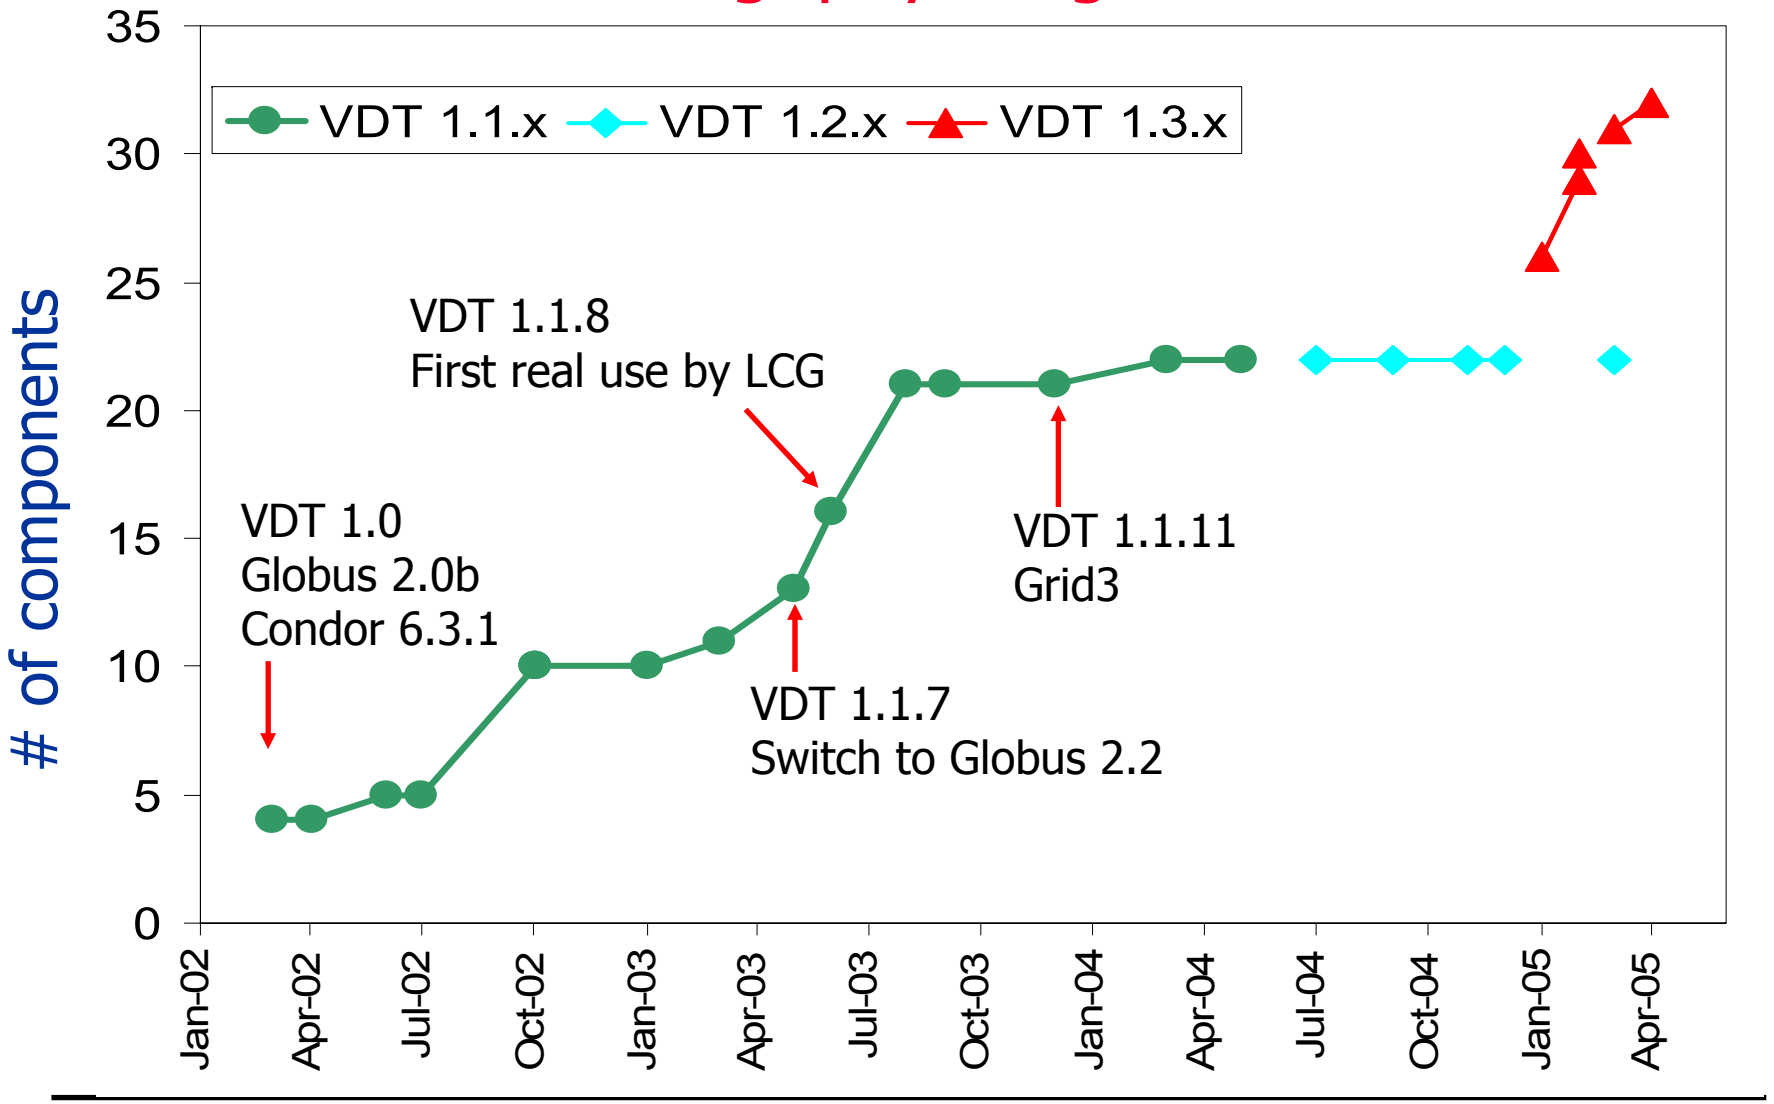

Trillium Science Drivers

- Experiments at Large Hadron Collider
  - New fundamental particles and forces
  - 100s of Petabytes 2007 - ?
- High Energy & Nuclear Physics expts
  - Top quark, nuclear matter at extreme density
  - ~1 Petabyte (1000 TB) 1997 – present
- LIGO (gravity wave search)
  - Search for gravitational waves
  - 100s of Terabytes 2002 – present
- Sloan Digital Sky Survey
  - Systematic survey of astronomical objects
  - 10s of Terabytes 2001 – present

LHC: Petascale Global Science

- Complexity: Millions of individual detector channels
- Scale: PetaOps (CPU), 100s of Petabytes (Data)
- Distribution: Global distribution of people & resources

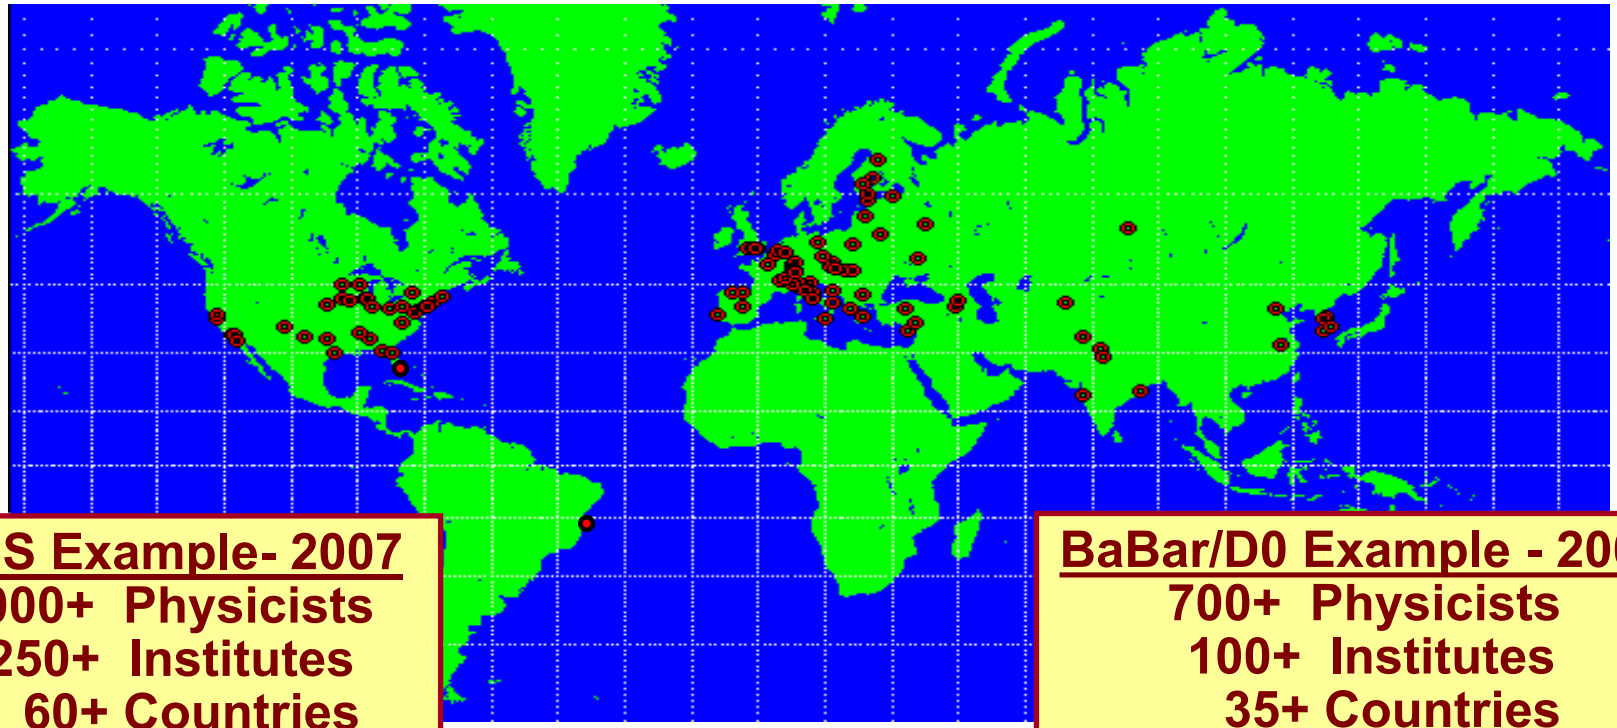

CMS Example- 2007
5000+ Physicists
250+ Institutes
60+ Countries

BaBar/D0 Example - 2004
700+ Physicists
100+ Institutes
35+ Countries

LHC Global Data Grid (2007+)

- 5000 physicists, 60 countries
- 10s of Petabytes/yr by 2008
- 1000 Petabytes in < 10 yrs?

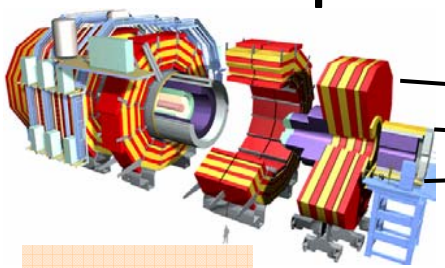

University LHC Tier2 Centers

- Tier2 facility
  - Essential university role in extended computing infrastructure
  - 20 – 25% of Tier1 national laboratory, supported by NSF
  - Validated by 3 years of experience (CMS, ATLAS)
- Functions
  - Perform physics analysis, simulations
  - Support experiment software, smaller institutions
- Official role in Grid hierarchy (U.S.)
  - Sanctioned by MOU with parent organization (ATLAS, CMS)
  - Selection by collaboration via careful process

Grids and Globally Distributed Teams

- **Non-hierarchical: Chaotic analyses + productions**
- **Superimpose significant random data flows**

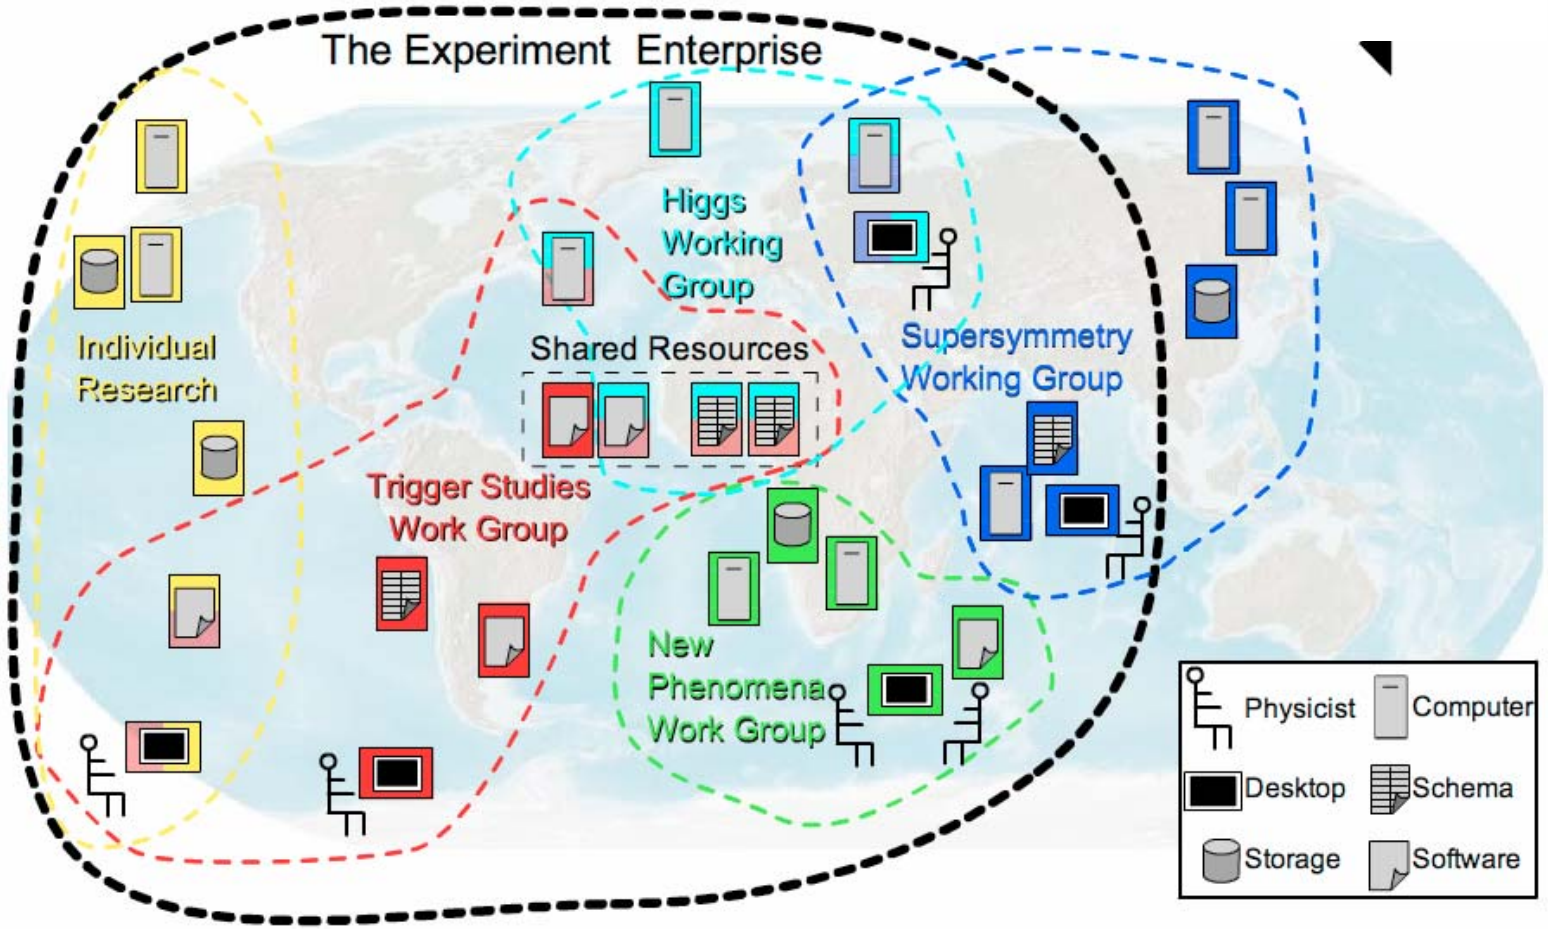

CMS: Grid Enabled Analysis Architecture

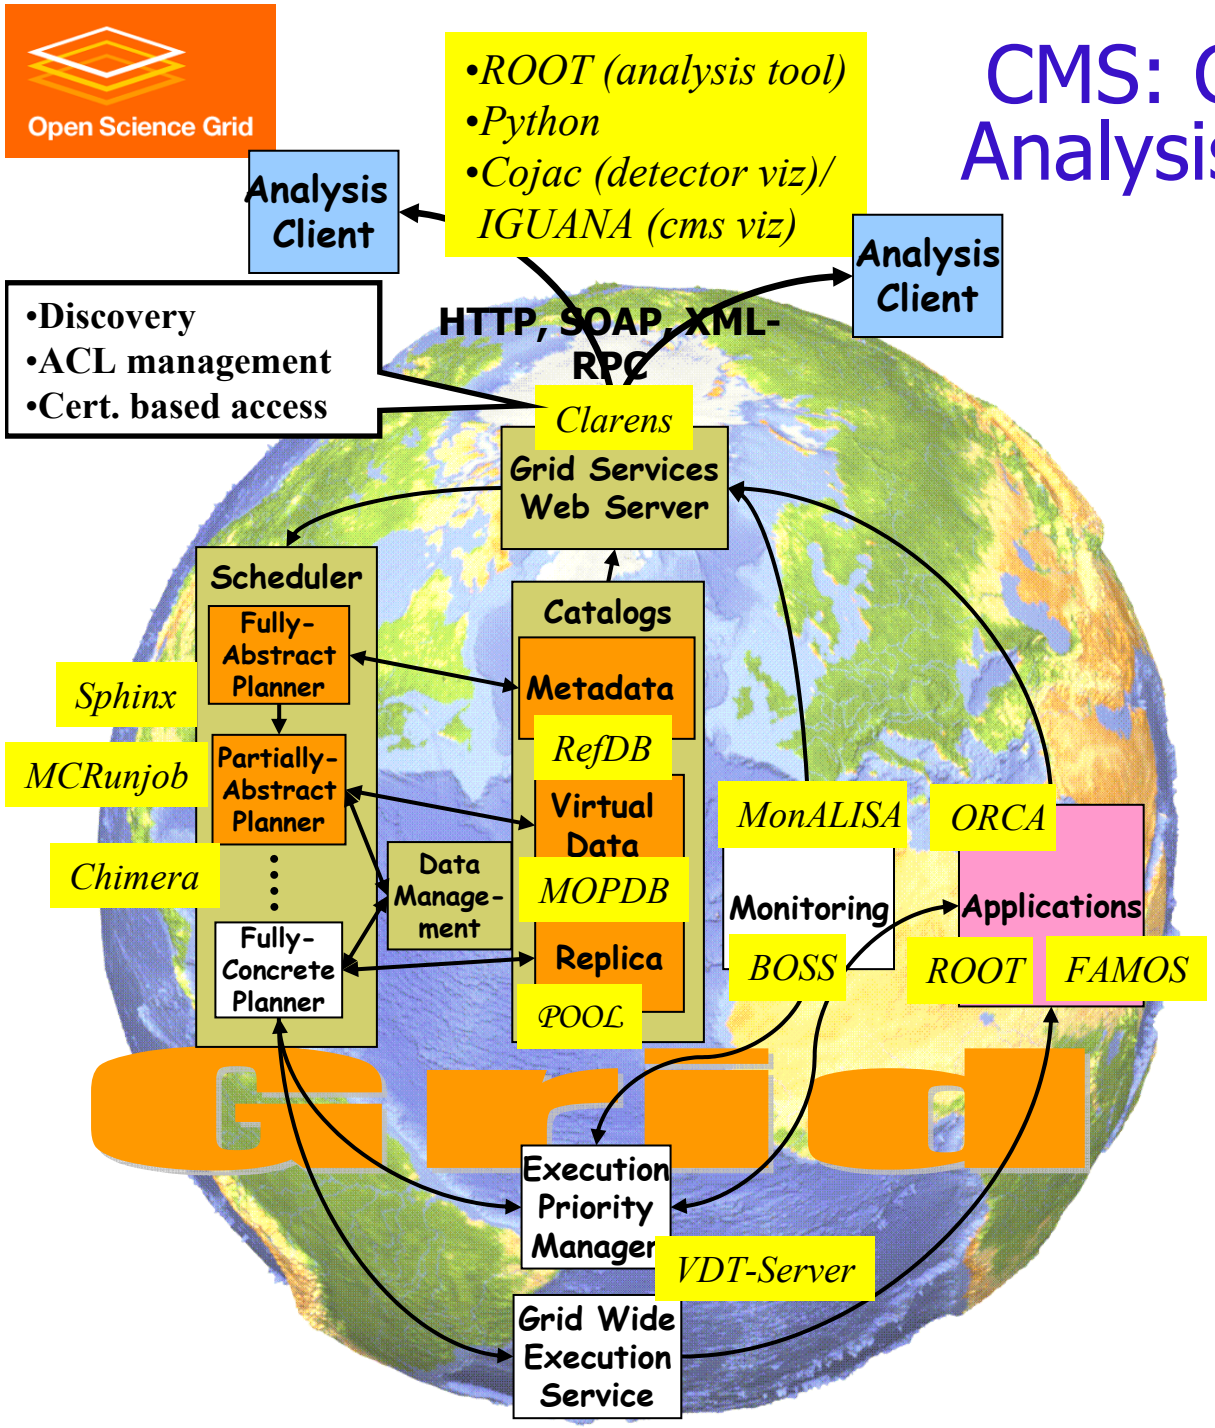

- Clients talk standard protocols to "Grid Services Web Server"
- Simple Web service API allows simple or complex analysis clients
- Typical clients: ROOT, Web Browser, ….
- Clarens portal hides complexity
- Key features: *Global* Scheduler, Catalogs, Monitoring, Grid-wide Execution service

Grid3: A National Grid Infrastructure

- 32 sites, 4000 CPUs: Universities + 4 national labs
- Part of LHC Grid, Running since October 2003
- Sites in US, Korea, Brazil, Taiwan
- Applications in HEP, LIGO, SDSS, Genomics, fMRI, CS

Grid3 Components

- Computers & storage at ~30 sites: 4000 CPUs
- Uniform service environment at each site
  - Globus 3.2: Authentication, execution management, data movement
  - Pacman: Installation of numerous VDT and application services
- Global & virtual organization services
  - Certification & reg. authorities, VO membership & monitor services
- Client-side tools for data access & analysis
  - Virtual data, execution planning, DAG management, execution management, monitoring
- IGOC: iVDGL Grid Operations Center
- Grid testbed: Grid3dev
  - Middleware development and testing, new VDT versions, etc.

Grid3 Applications

CMS experiment	p-p collision simulations & analysis
ATLAS experiment	p-p collision simulations & analysis
BTEV experiment	p-p collision simulations & analysis
LIGO	Search for gravitational wave sources
SDSS	Galaxy cluster finding
Bio-molecular analysis	Shake n Bake (SnB) (Buffalo)
Genome analysis	GADU/Gnare
fMRI	Functional MRI (Dartmouth)
CS Demonstrators	Job Exerciser, GridFTP, NetLogger

www.ivdgl.org/grid3/applications

Grid3 Shared Use Over 6 months

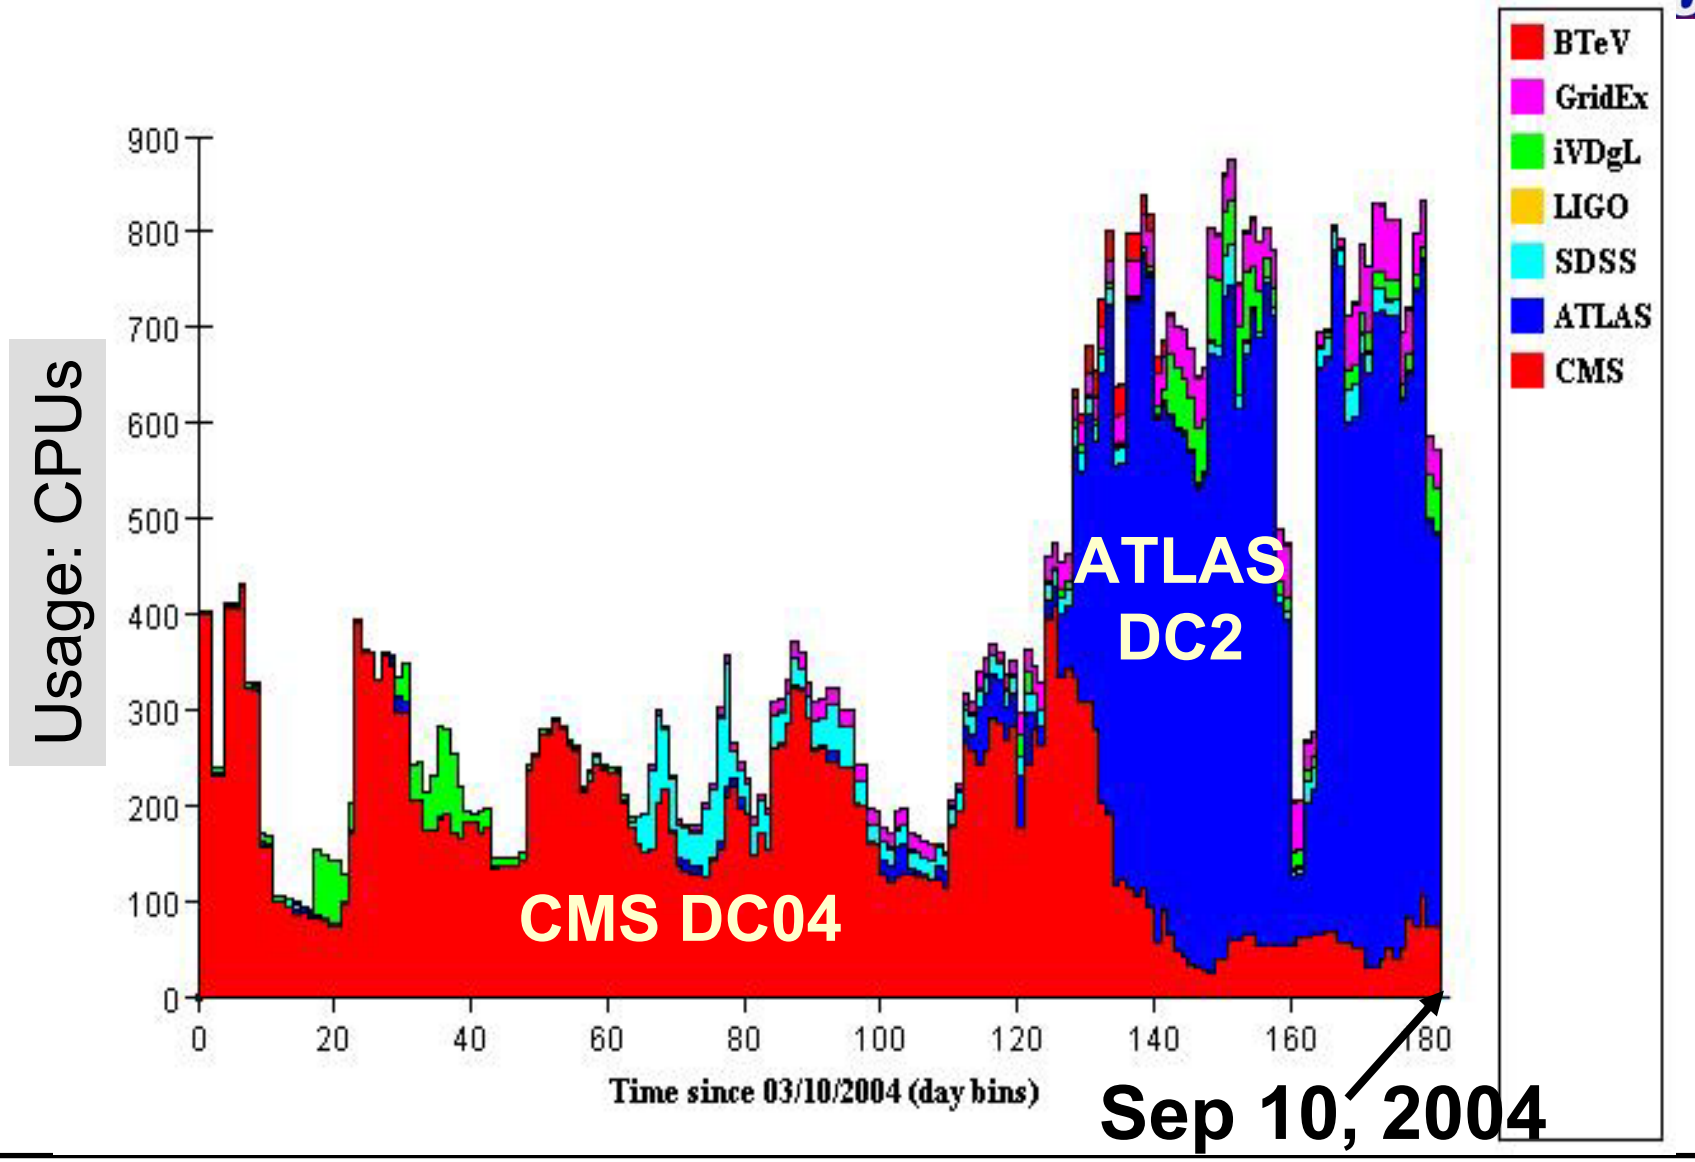

Grid3 Production Over 13 Months

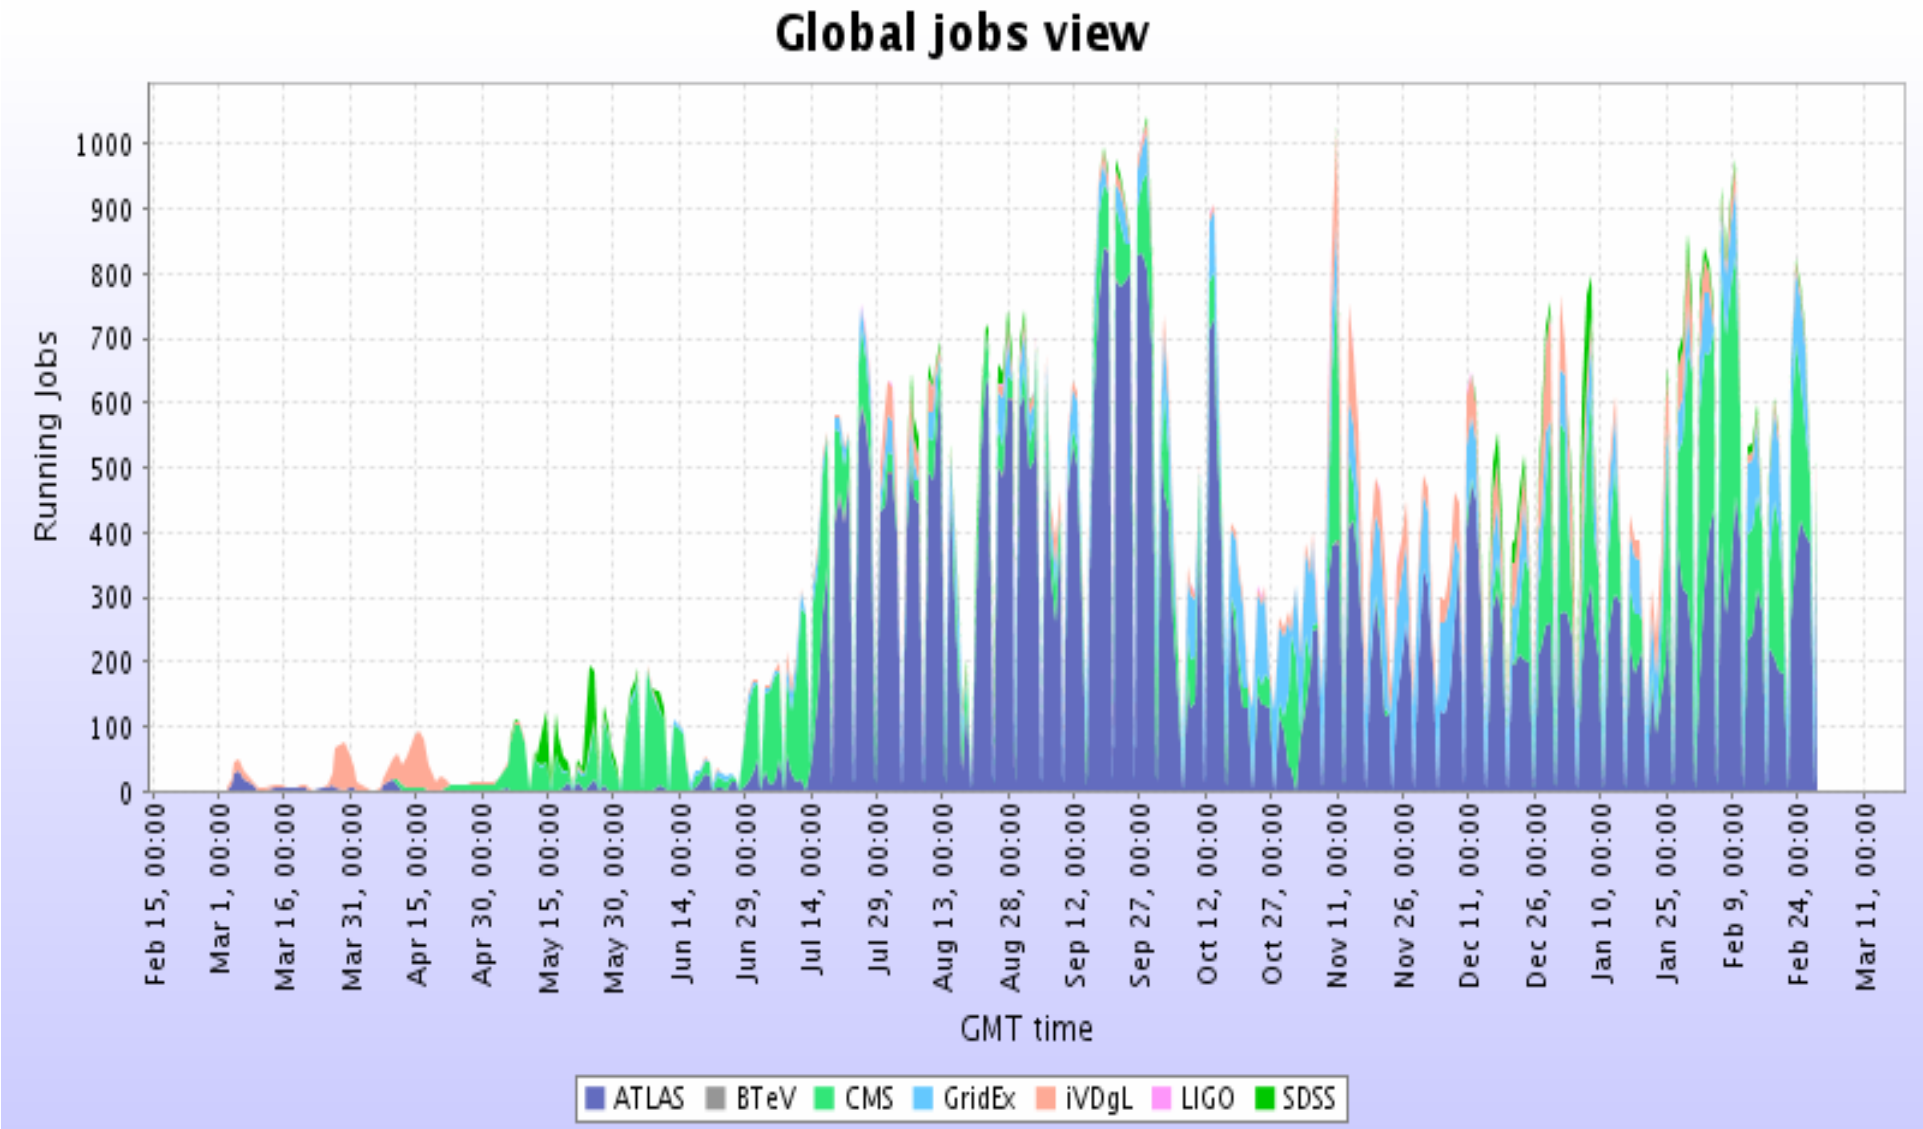

U.S. CMS 2003 Production

- 10M p-p collisions; largest ever
  - 2× simulation sample
  - ½ manpower
- Multi-VO sharing

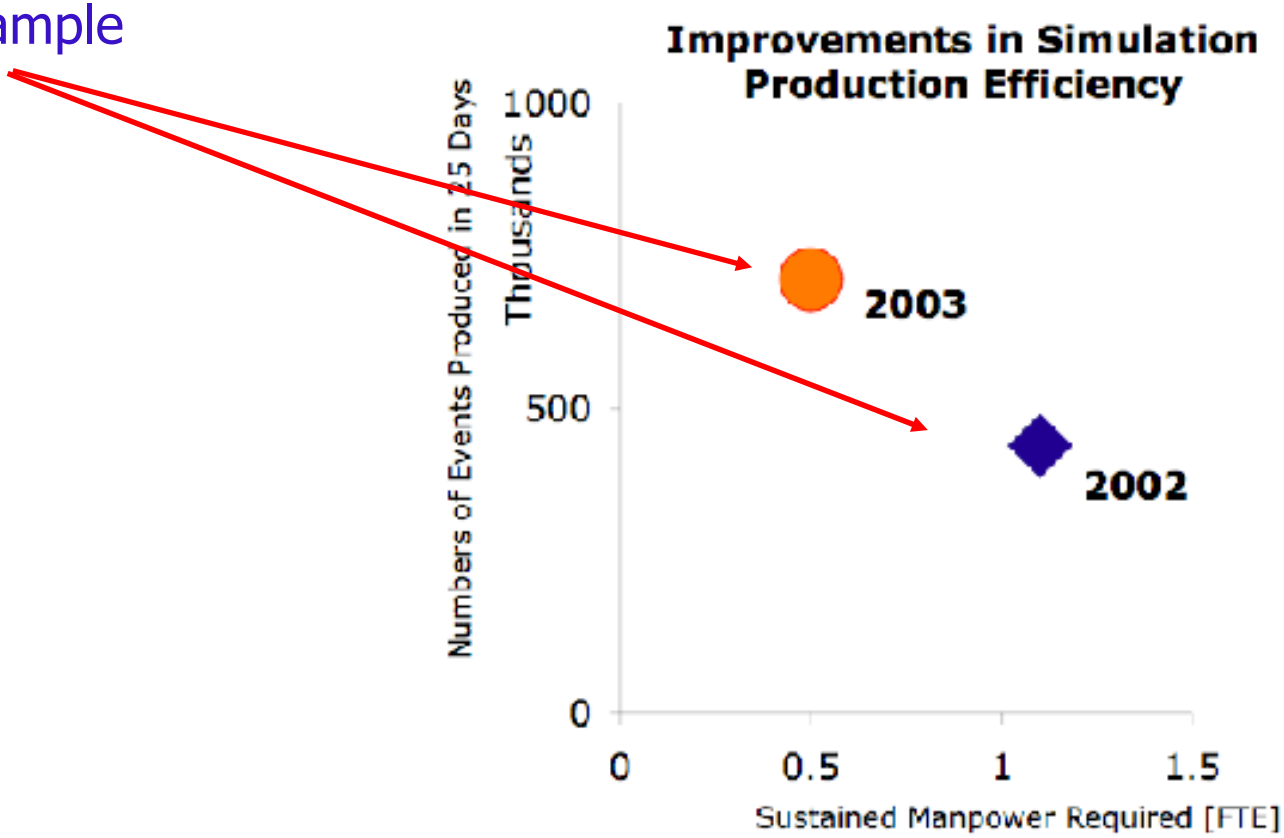

Grid3 Lessons Learned

- How to **operate** a Grid as a facility
  - Tools, services, error recovery, procedures, docs, organization
  - Delegation of responsibilities (Project, VO, service, site, …)
  - Crucial role of Grid Operations Center (GOC)
- How to **support** people ⇔ people relations
  - Face-face meetings, phone cons, 1-1 interactions, mail lists, etc.
- How to **test** and **validate** Grid tools and applications
  - Vital role of testbeds
- How to **scale** algorithms, software, process
  - Some successes, but "interesting" failure modes still occur
- How to **apply** distributed cyberinfrastructure
  - Successful production runs for several applications

Grid3 ⇒ Open Science Grid

- Iteratively build & extend Grid3
  - Grid3 → OSG-0 → OSG-1 → OSG-2 → …
  - Shared resources, benefiting broad set of disciplines
  - Grid middleware based on Virtual Data Toolkit (VDT)
- Consolidate elements of OSG collaboration
  - Computer and application scientists
  - Facility, technology and resource providers (labs, universities)
- Further develop OSG
  - Partnerships with other sciences, universities
  - Incorporation of advanced networking
  - Focus on general services, operations, end-to-end performance
- Aim for July 2005 deployment

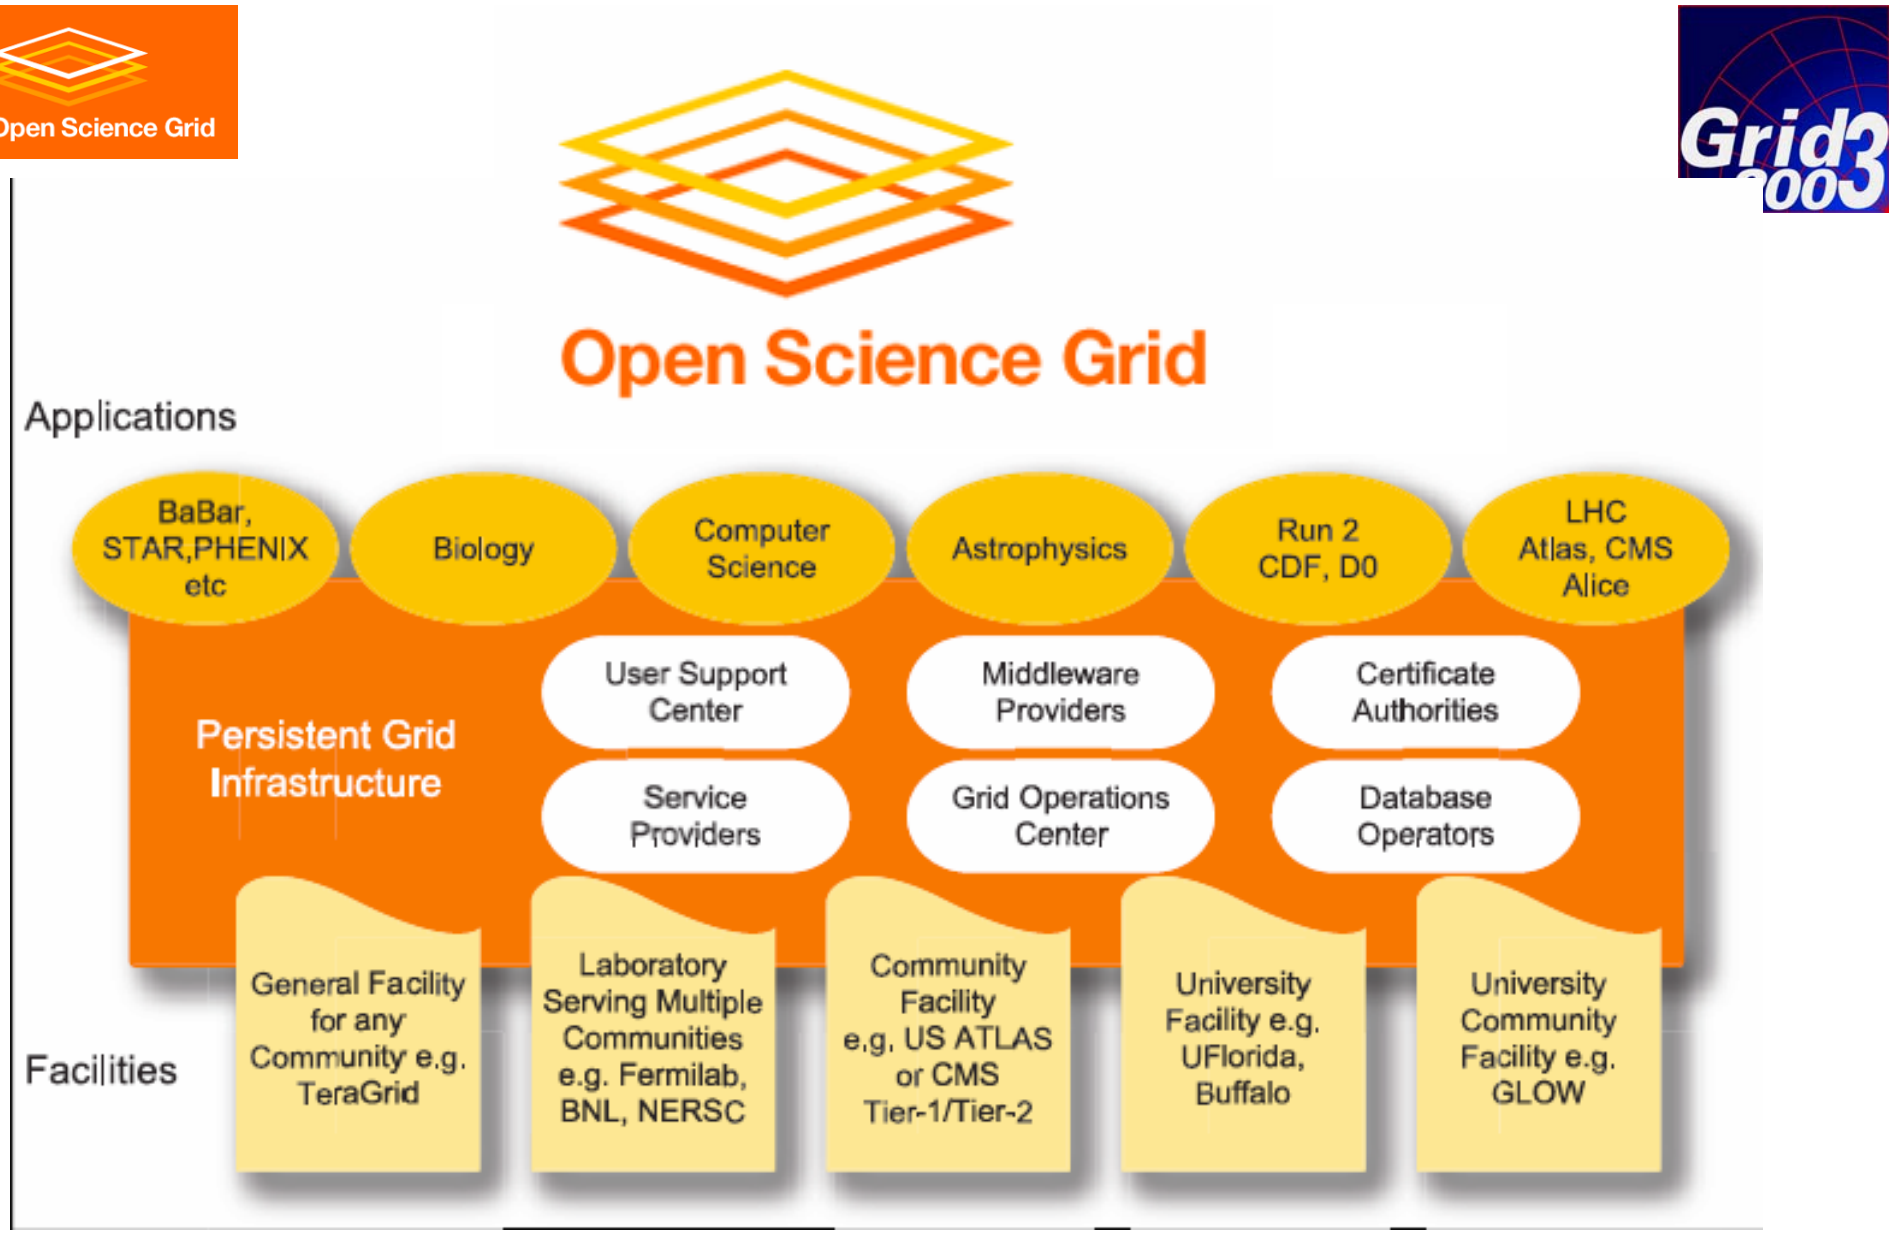

http://www.opensciencegrid.org

OSG Organization

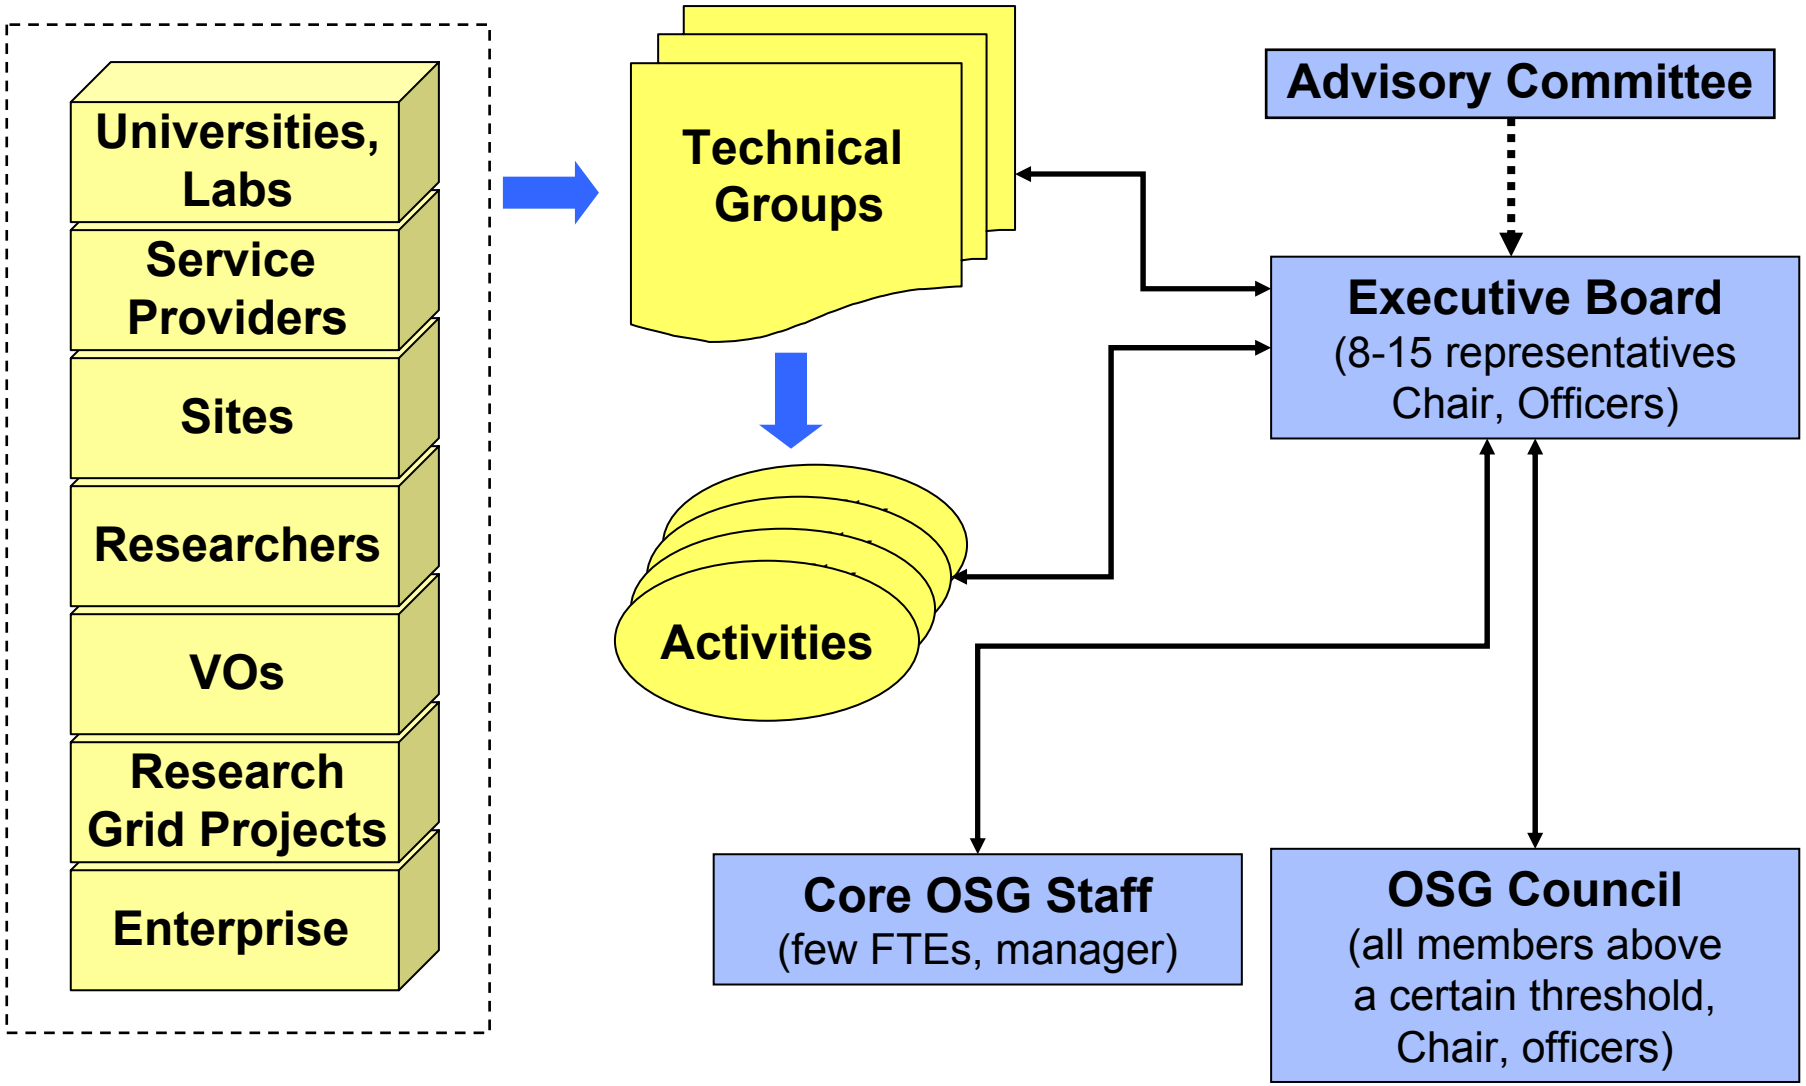

OSG Technical Groups & Activities

- Technical Groups address and coordinate technical areas
  - Propose and carry out activities related to their given areas
  - Liaise & collaborate with other peer projects (U.S. & international)
  - Participate in relevant standards organizations.
  - Chairs participate in Blueprint, Integration and Deployment activities
- Activities are well-defined, scoped tasks contributing to OSG
  - Each Activity has deliverables and a plan
  - … is self-organized and operated
  - … is overseen & sponsored by one or more Technical Groups

TGs and Activities are where the real work gets done

OSG Technical Groups

Governance	Charter, organization, by-laws, agreements, formal processes
Policy	VO & site policy, authorization, priorities, privilege & access rights
Security	Common security principles, security infrastructure
Monitoring and Information Services	Resource monitoring, information services, auditing, troubleshooting
Storage	Storage services at remote sites, interfaces, interoperability
Support Centers	Infrastructure and services for user support, helpdesk, trouble ticket
Education / Outreach	Training, interface with various E/O projects
Networks (new)	Including interfacing with various networking projects

OSG Activities

Blueprint	Defining principles and best practices for OSG
Deployment	Deployment of resources & services
Provisioning	Connected to deployment
Incidence response	Plans and procedures for responding to security incidents
Integration	Testing & validating & integrating new services and technologies
Data Resource Management (DRM)	Deployment of specific Storage Resource Management technology
Documentation	Organizing the documentation infrastructure
Accounting	Accounting and auditing use of OSG resources
Interoperability	Primarily interoperability between
Operations	Operating Grid-wide services

The Path to the OSG Operating Grid

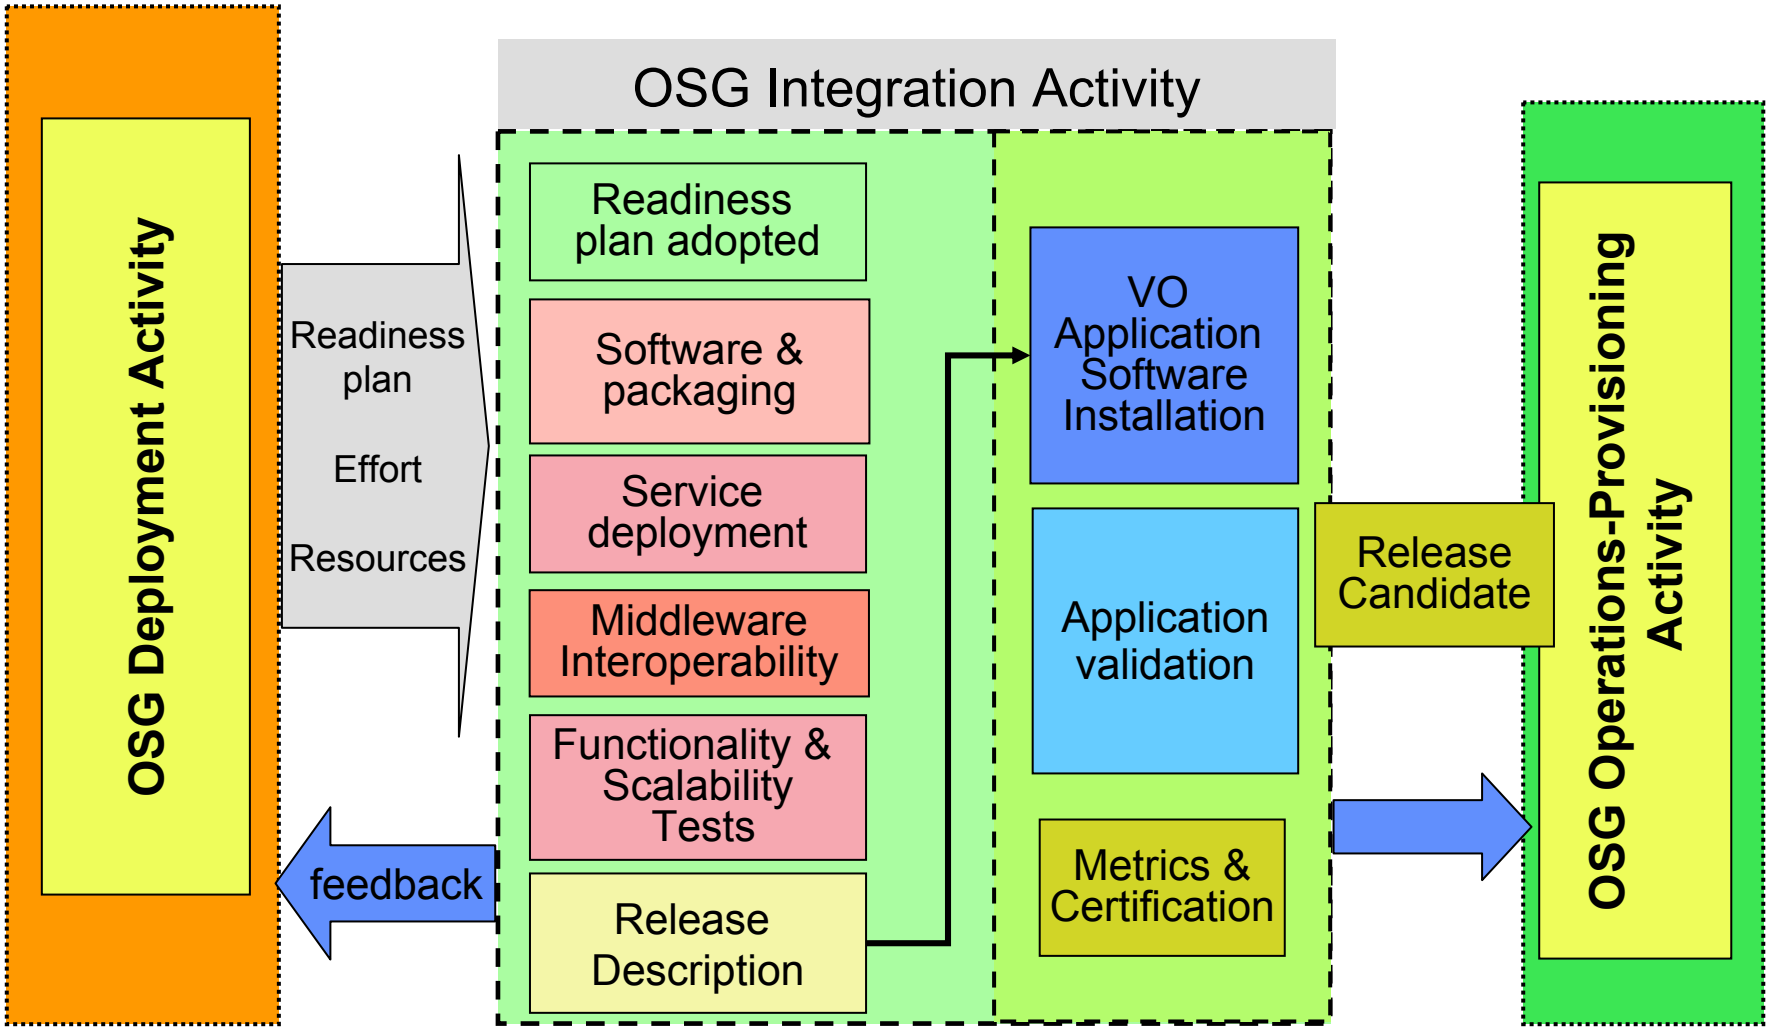

OSG Integration Testbed
>20 Sites and Rising

Status of OSG Deployment

- OSG infrastructure release "accepted" for deployment
  - US CMS application "flood testing" successful
  - D0 simulation & reprocessing jobs running on selected OSG sites
  - Others in various stages of readying applications & infrastructure (ATLAS, CMS, STAR, CDF, BaBar, fMRI)

- Deployment process underway: End of July?
  - Open OSG and transition resources from Grid3
  - Applications will use growing ITB & OSG resources during transition

http://osg.ivdgl.org/twiki/bin/view/Integration/WebHome

Connections to LCG and EGEE

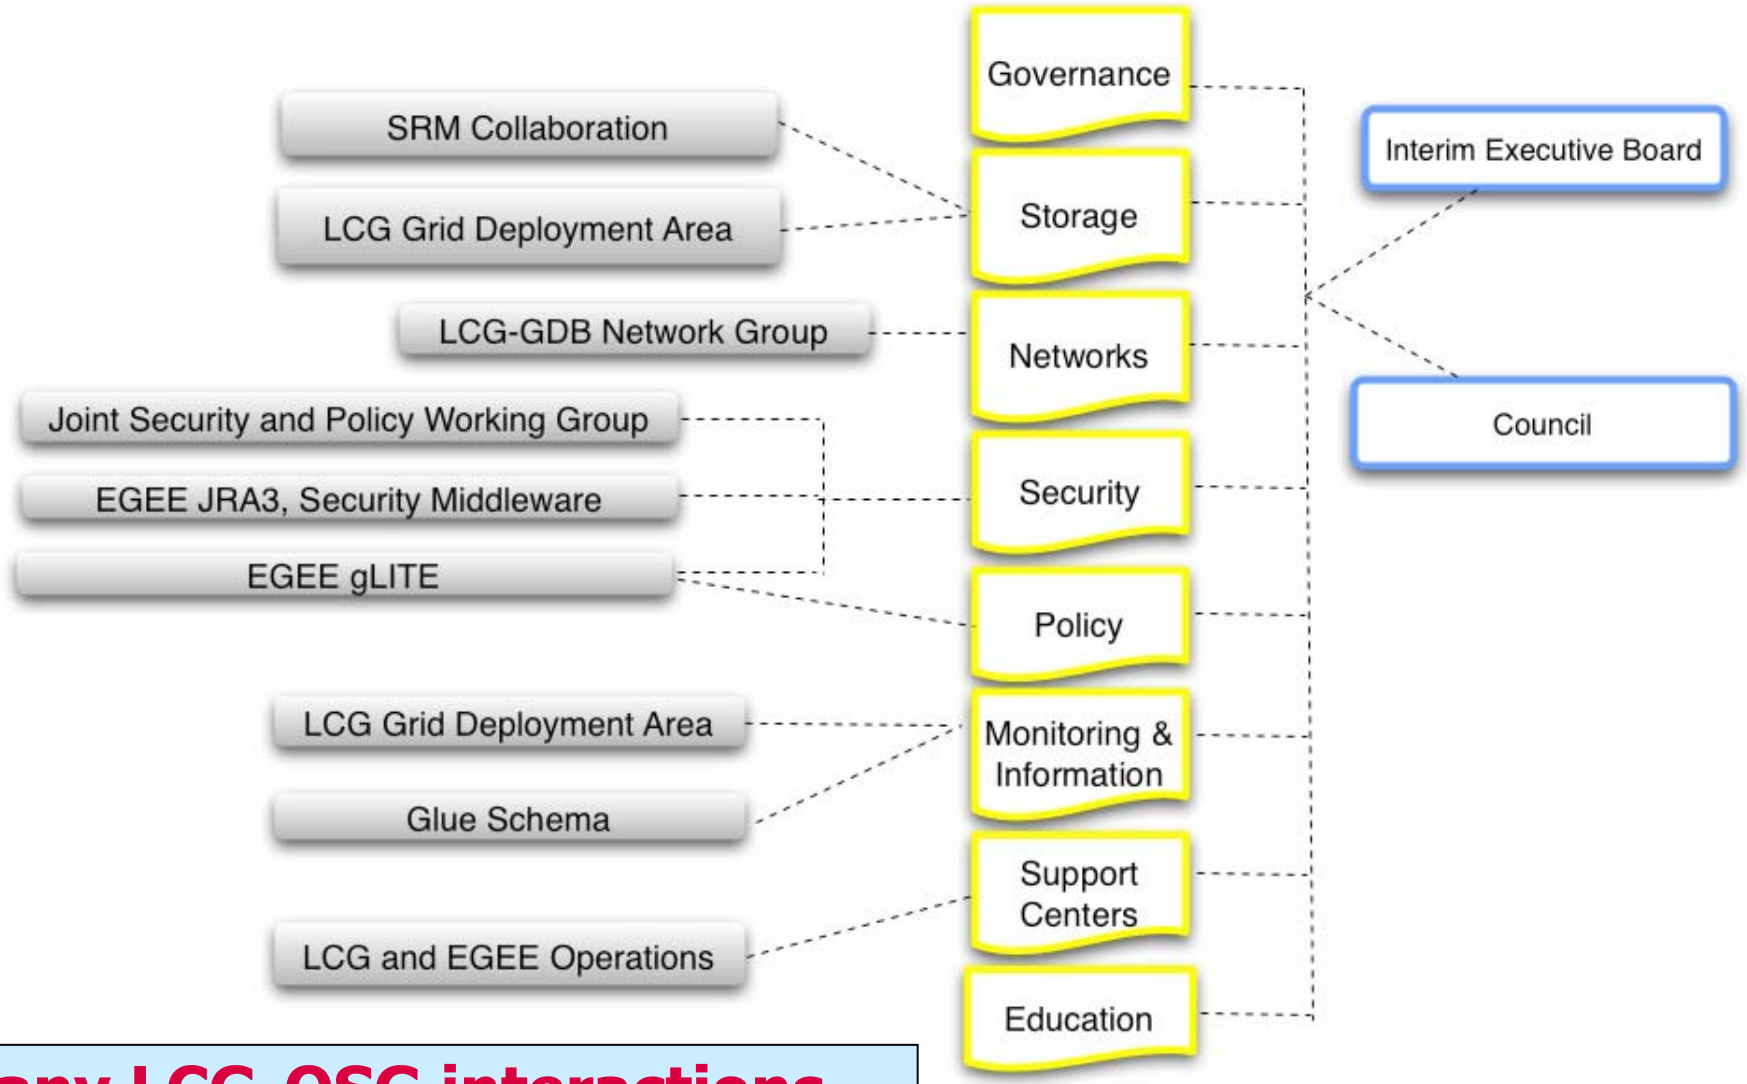

Many LCG-OSG interactions

Interoperability & Federation

- Transparent use of federated Grid infrastructures a goal
  - LCG, EGEE
  - TeraGrid
  - State-wide Grids
  - Campus Grids (Wisconsin, Florida, etc)
- Some early activities with LCG
  - Some OSG/Grid3 sites appear in LCG map
  - D0 bringing reprocessing to LCG sites through adaptor node
  - CMS and ATLAS can run their jobs on both LCG and OSG
- Increasing interaction with TeraGrid
  - CMS and ATLAS sample simulation jobs are running on TeraGrid
  - Plans for TeraGrid allocation for jobs running in Grid3 model (group accounts, binary distributions, external data management, etc)

UltraLight: Advanced Networking in Applications

Funded by ITR2004

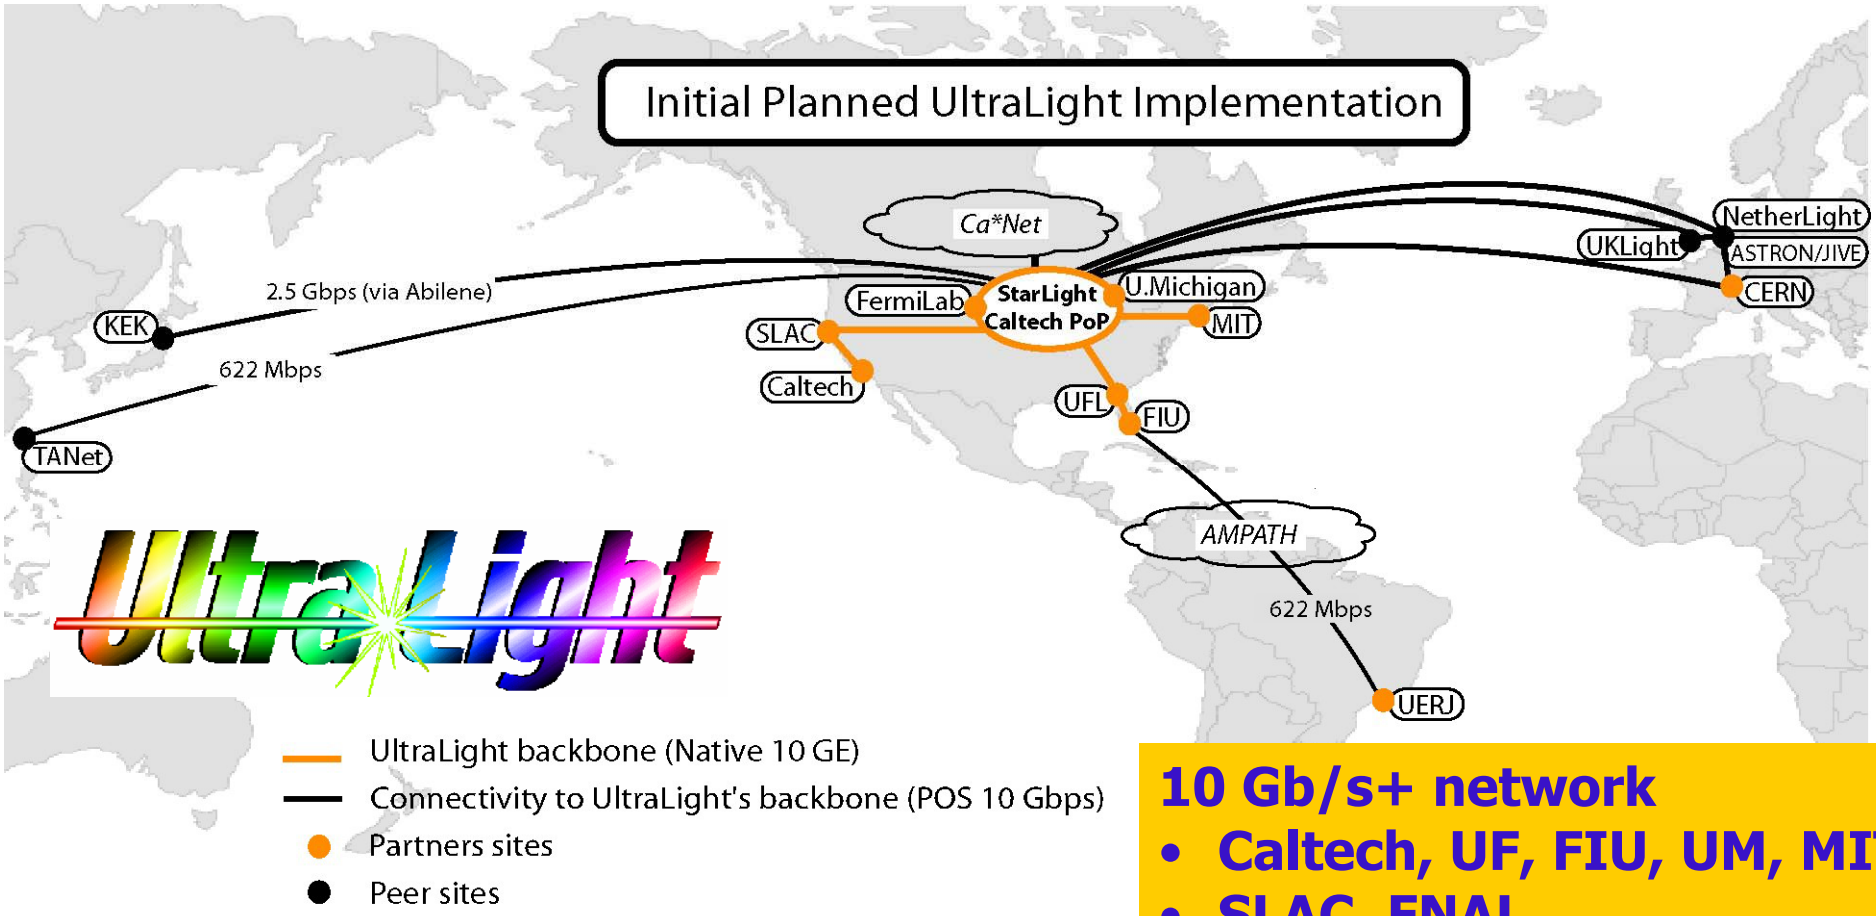

10 Gb/s+ network
- **Caltech, UF, FIU, UM, MIT**
- **SLAC, FNAL**
- **Int'l partners**
- **Level(3), Cisco, NLR**

UltraLight: New Information System

- A new class of integrated information systems
  - Includes networking as a managed resource for the first time
  - Uses "Hybrid" packet-switched and circuit-switched optical network infrastructure
  - Monitor, manage & optimize network and Grid Systems in realtime
- Flagship applications: HEP, eVLBI, "burst" imaging
  - "Terabyte-scale" data transactions in minutes
  - Extend Real-Time eVLBI to the 10 – 100 Gb/s Range
- Powerful testbed
  - Significant storage, optical networks for testing new Grid services
- Strong vendor partnerships
  - Cisco, Calient, NLR, CENIC, Internet2/Abilene

iVDGL, GriPhyN Education/Outreach

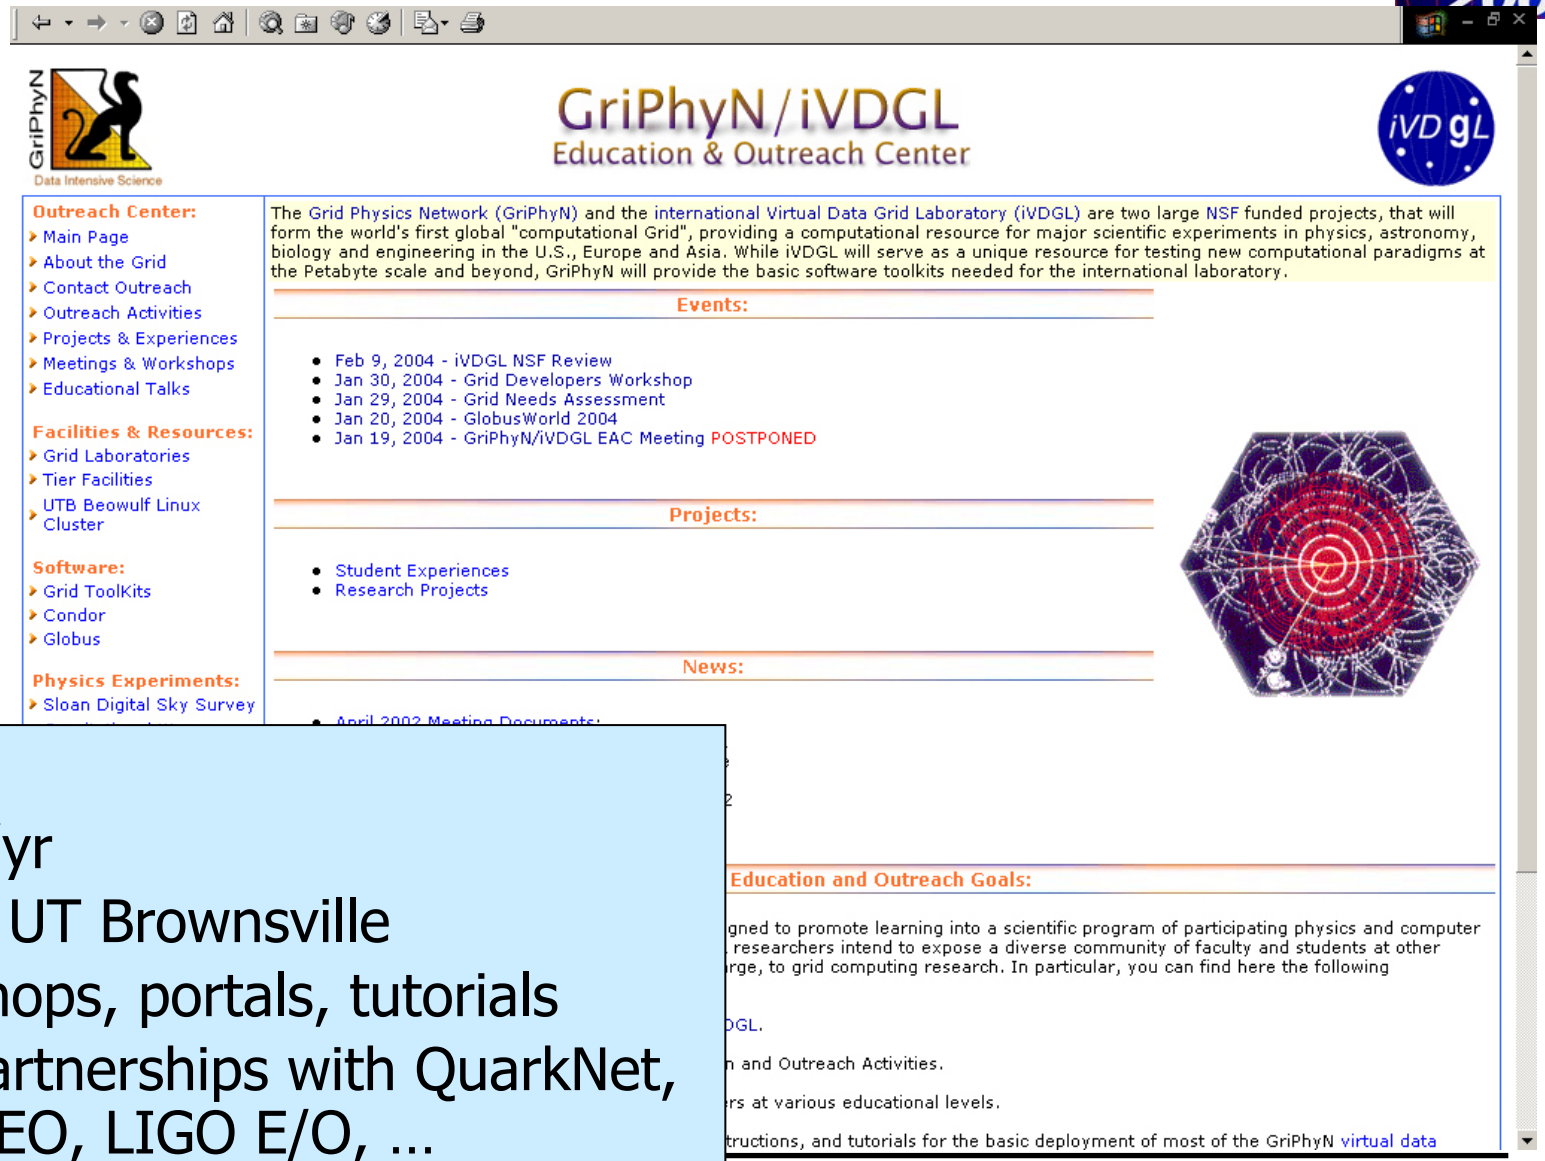

Basics

- $200K/yr
- Led by UT Brownsville
- Workshops, portals, tutorials
- New partnerships with QuarkNet, CHEPREO, LIGO E/O, …

U.S. Grid Summer School

- First of its kind in the U.S. (June 2004, South Padre Island)
  - 36 students, diverse origins and types (M, F, MSIs, etc)
- Marks new direction for U.S. Grid efforts
  - First attempt to systematically train people in Grid technologies
  - First attempt to gather relevant materials in one place
  - Today: Students in CS and Physics
  - Next: Students, postdocs, junior & senior scientists
- Reaching a wider audience
  - Put lectures, exercises, video, on the web
  - More tutorials, perhaps 2-3/year
  - Dedicated resources for remote tutorials
  - Create "Grid Cookbook", e.g. Georgia Tech
- Second workshop: July 11–15, 2005
  - South Padre Island

QuarkNet/GriPhyN e-Lab Project

Join a national collaboration of high school students to study cosmic rays.

http://quarknet.uchicago.edu/elab/cosmic/home.jsp

Spending all your time in a shower?

When you're sleeping or sitting in class, cosmic rays shower the earth and everything on it.

What are cosmic rays?

Where do they come from?

Where do they hit?

Some cosmic rays have so much energy that scientists are not sure where they come from. A number of reseach projects are looking at this question.

Password:

Login

Need login info?
Ask your teacher.

Working on your own?
Contact quarknet@fnal.gov.

To explore our website,

Who are we?

We're a collaboration of high school students and teach
cosmic ray data to answer some of these questions. W
scientists to provide cutting edge tools that use **grid tec**
graphs, and posters and collaborate with other students

Who can join?

You! Think about steps you'd take to investigate cosmi
started? What do you need to know? Can you collect ar

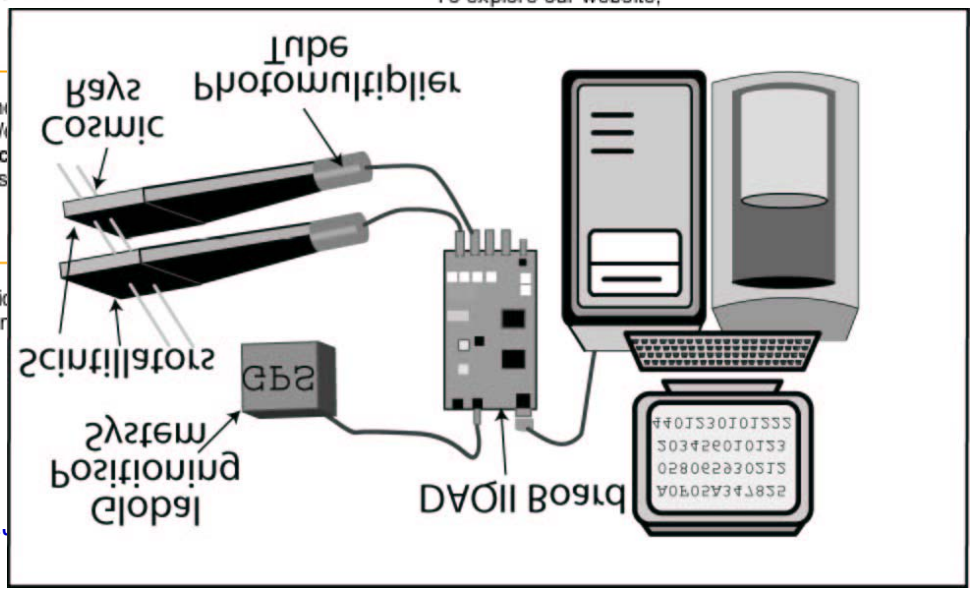

Student Muon Lifetime Analysis in GriPhyN/QuarkNet

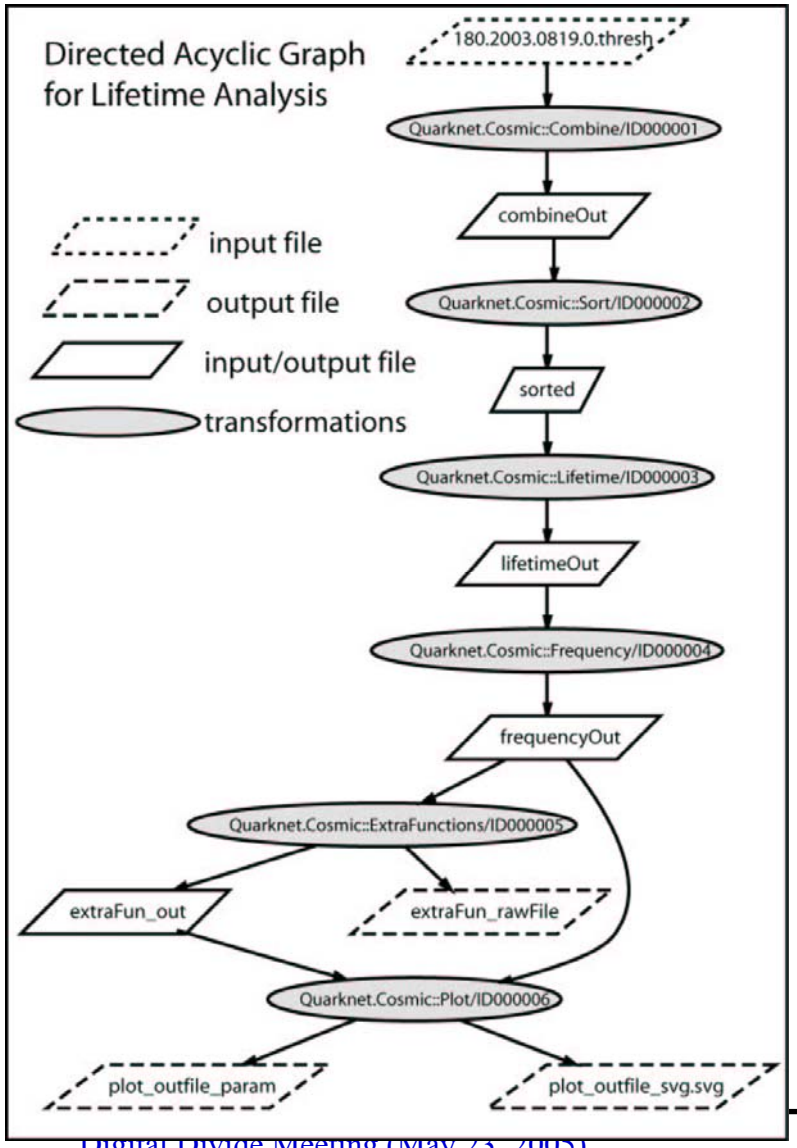

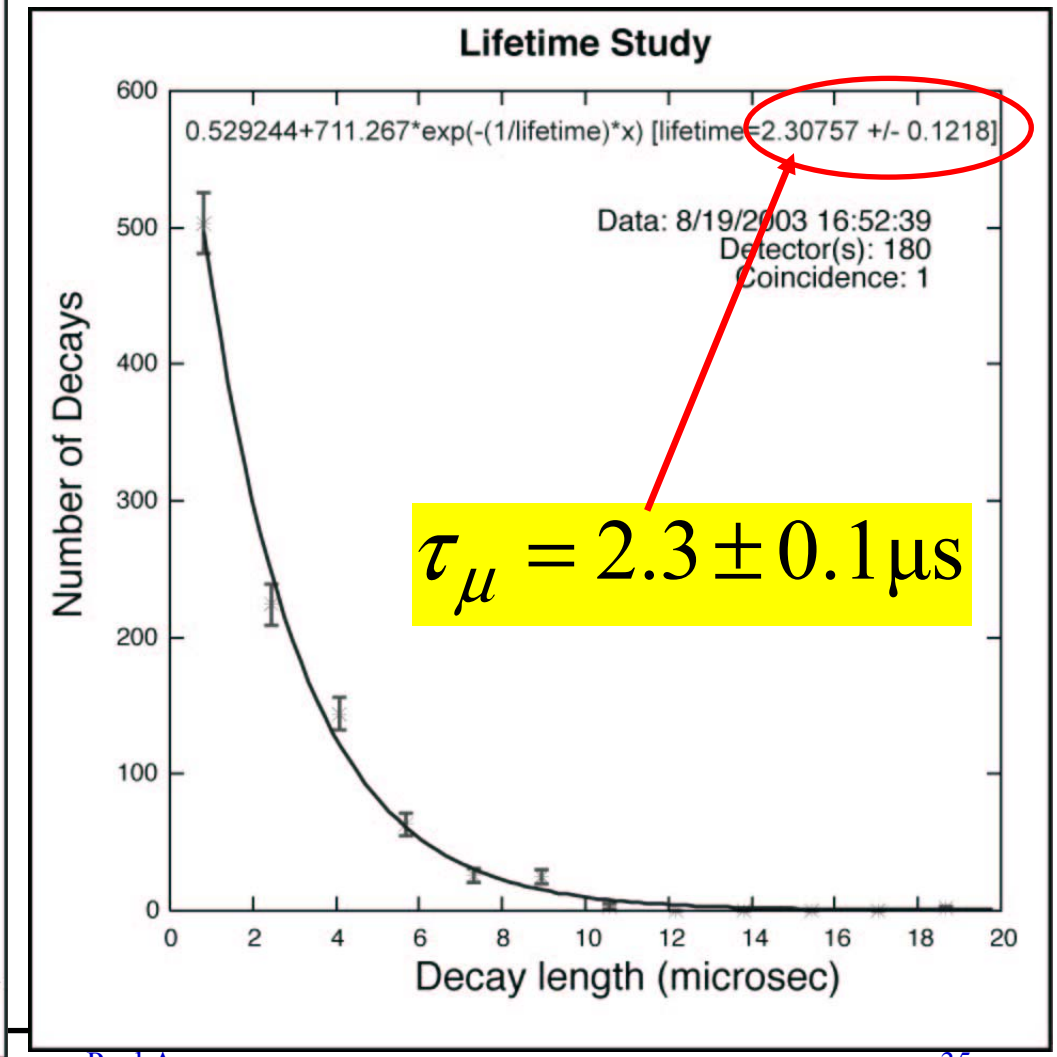

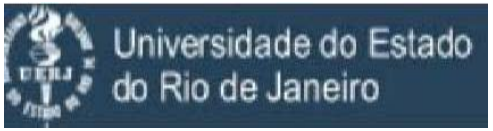

CHEPREO: Center for High Energy Physics Research and Educational Outreach

Florida International University

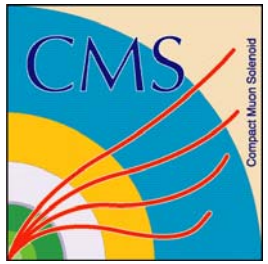

- Physics Learning Center
- CMS Research
- iVDGL Grid Activities
- AMPATH network (S. America)

- Funded September 2003
- $4M initially (3 years)
- MPS, CISE, EHR, INT

Science Grid Communications

Broad set of activities

- News releases, PR, etc.
- Science Grid This Week
- Katie Yurkewicz talk

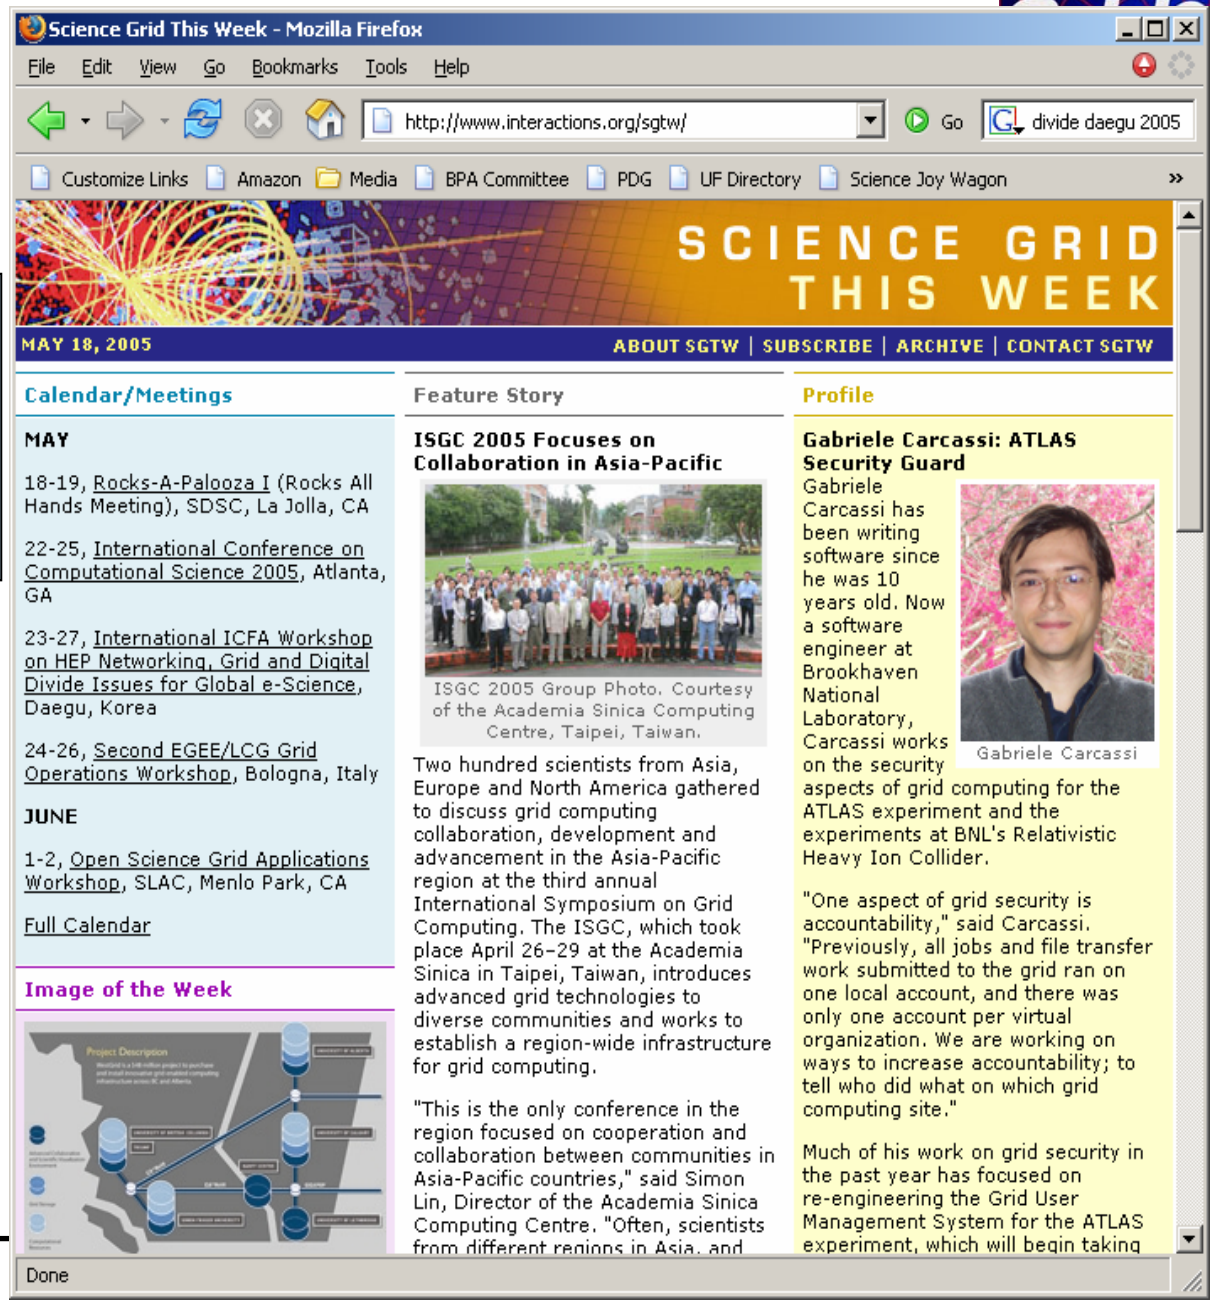

Grids and the Digital Divide
Rio de Janeiro + Daegu

HEPGRID AND DIGITAL DIVIDE WORKSHOP

Background

- World Summit on Information Society
- HEP Standing Committee on Inter-regional Connectivity (SCIC)

Themes

- Global collaborations, Grids and addressing the Digital Divide
- Focus on poorly connected regions
- Brazil (2004), Korea (2005)

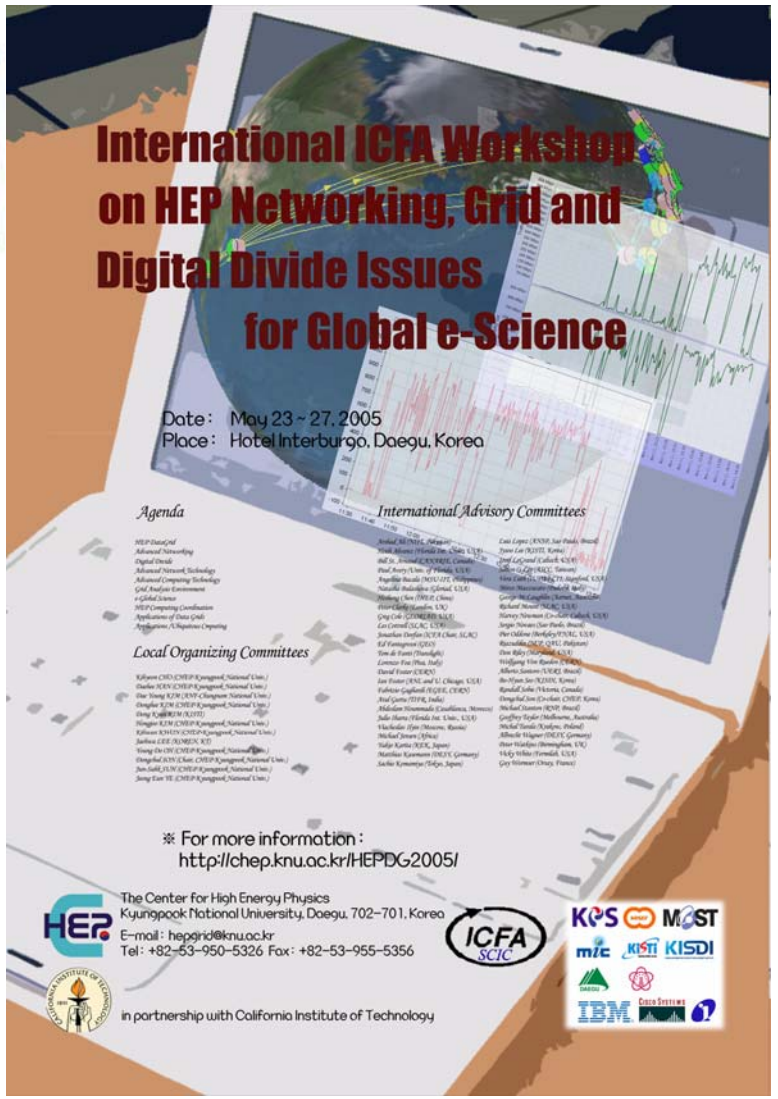

New Campus Research Grids (e.g., Florida)

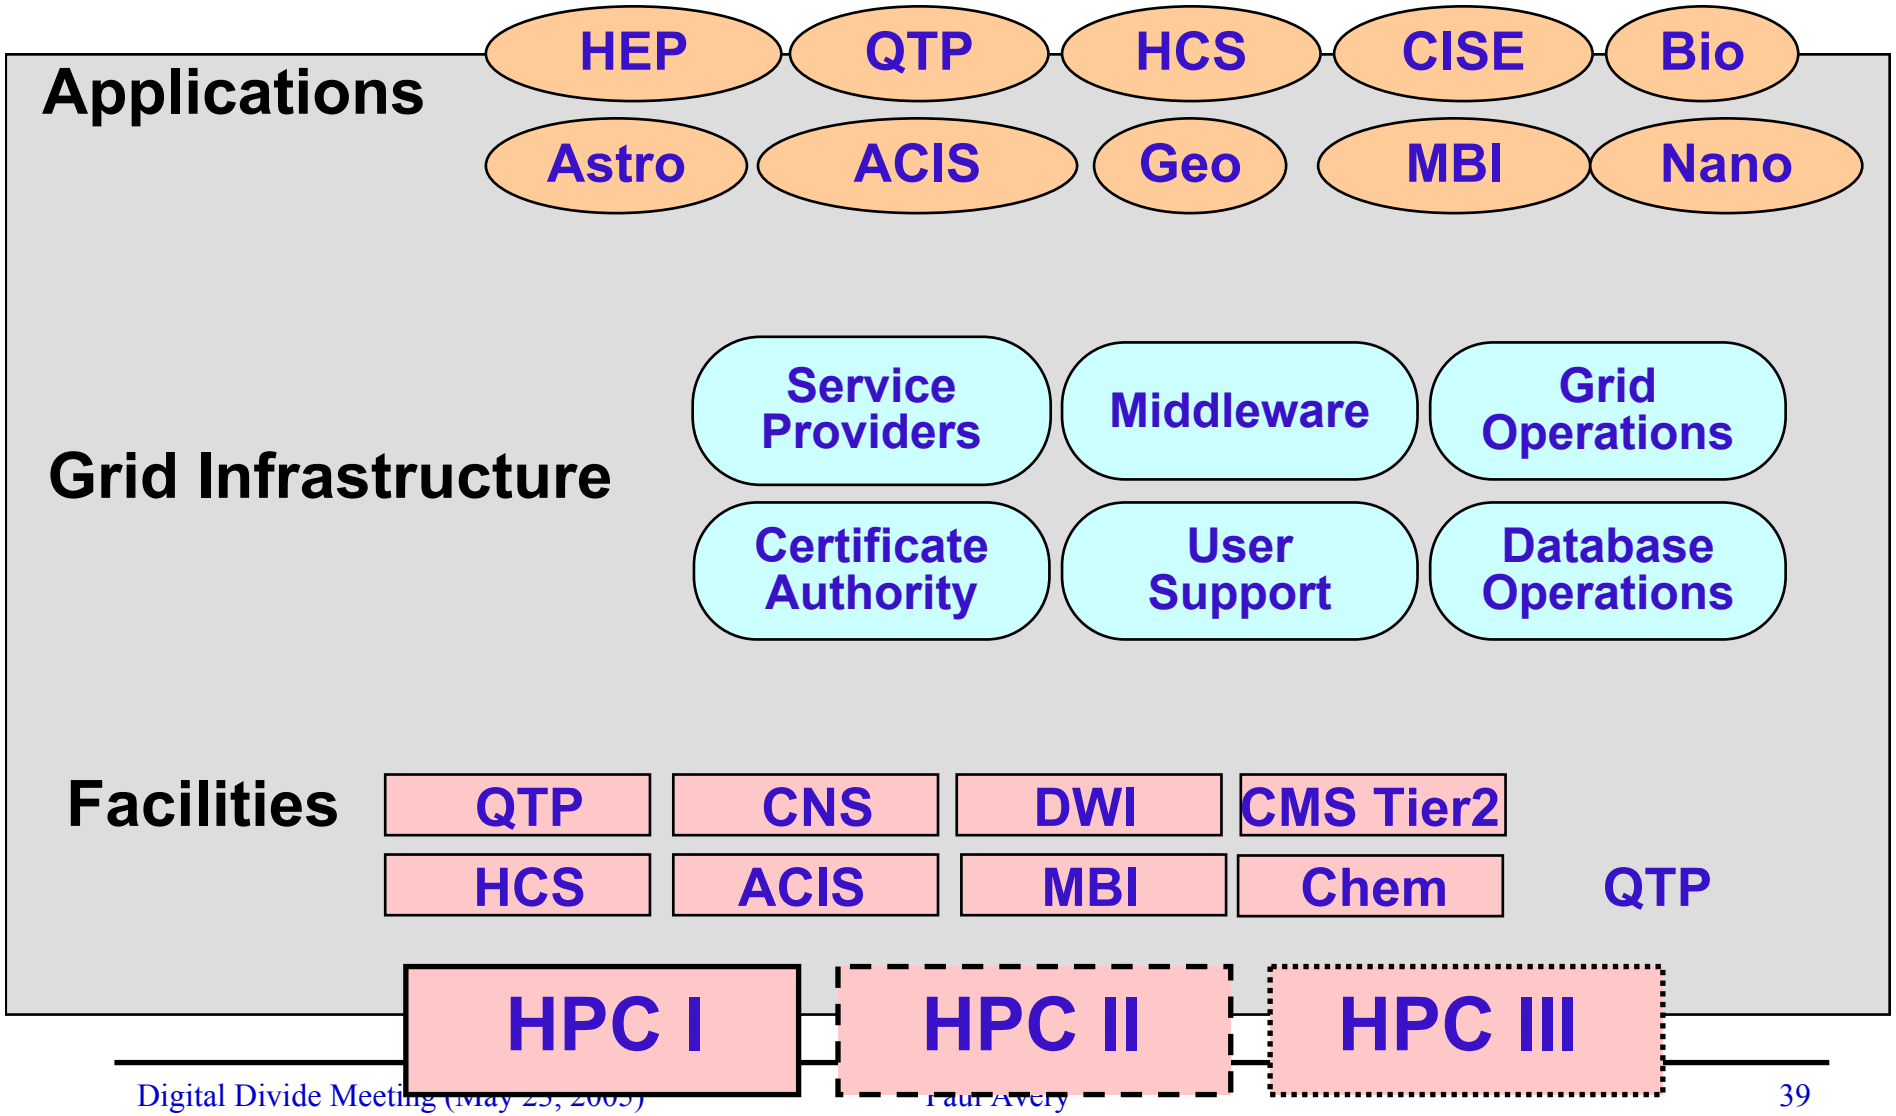

US HEP Data Grid Timeline

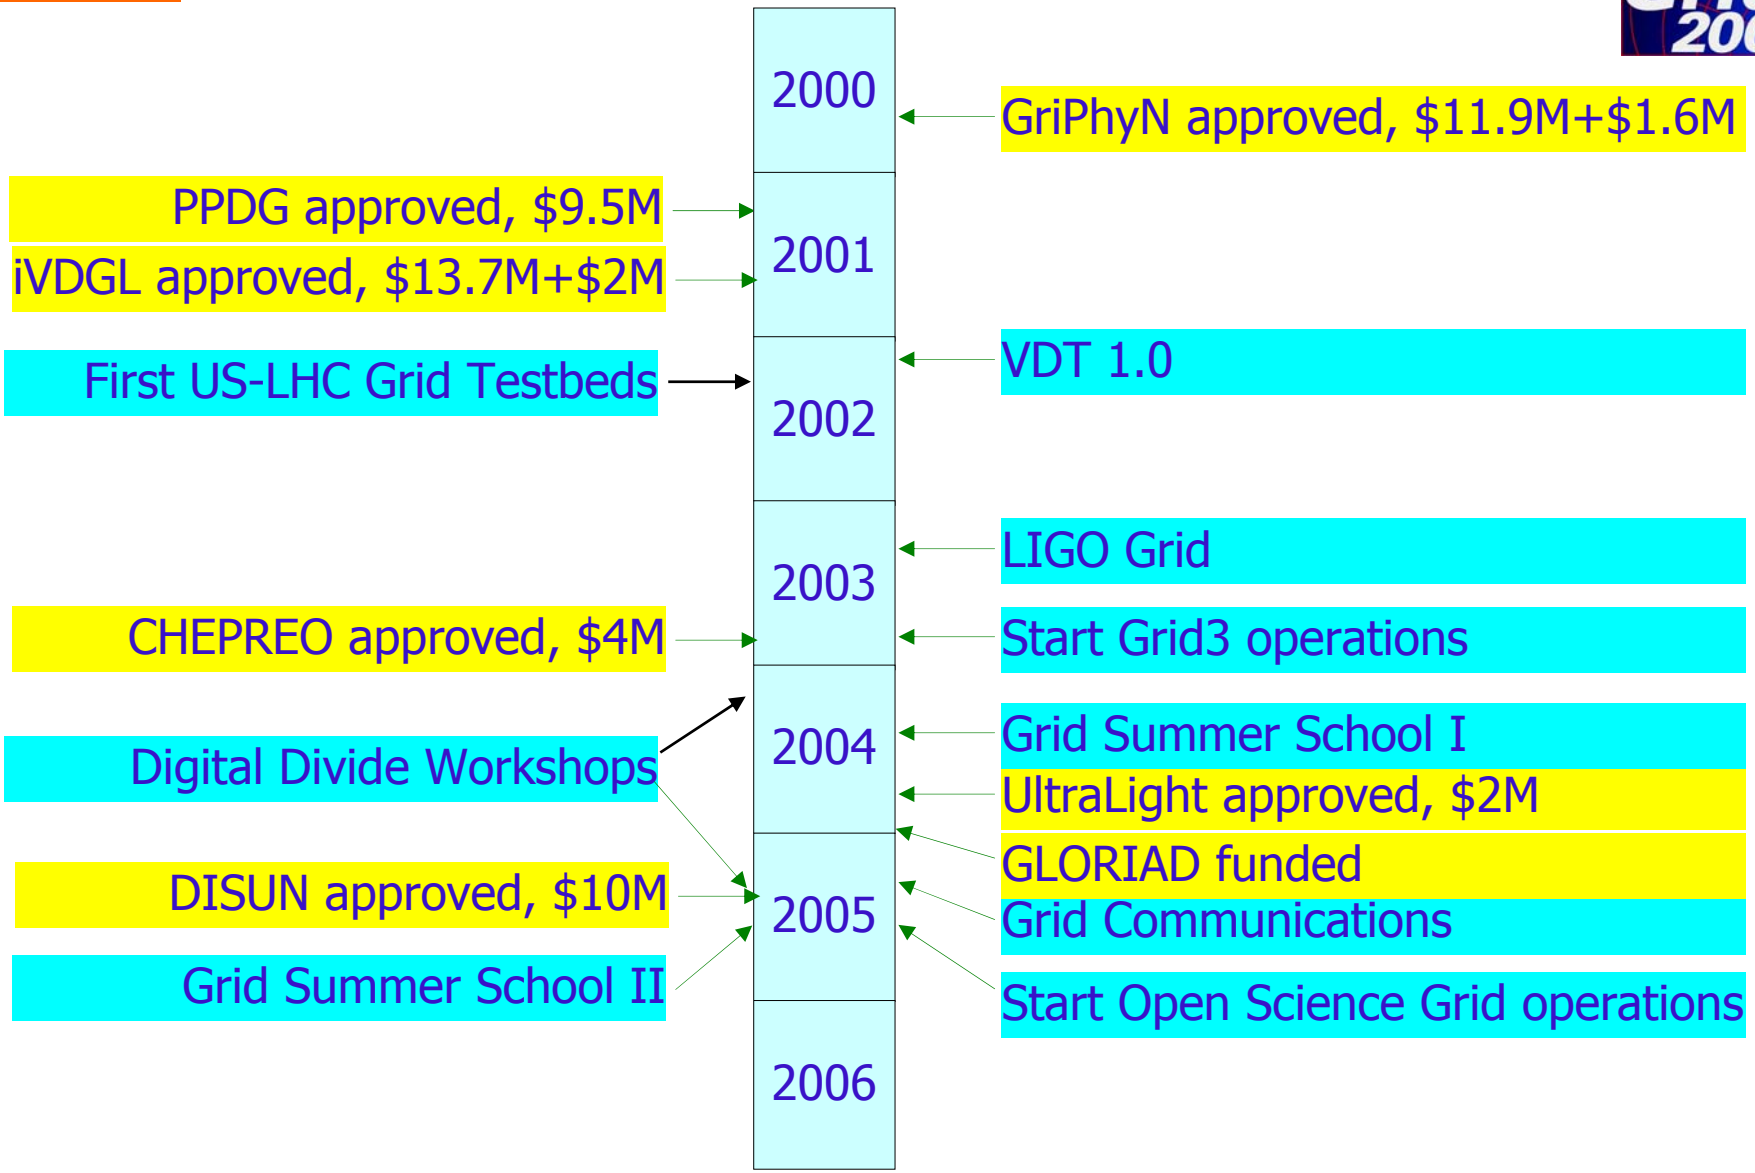

Summary

- Grids enable 21st century collaborative science
  - Linking research communities and resources for scientific discovery
  - Needed by global collaborations pursuing "petascale" science
- Grid3 was an important first step in developing US Grids
  - Value of planning, coordination, testbeds, rapid feedback
  - Value of learning how to operate a Grid as a facility
  - Value of building & sustaining community relationships
- Grids drive need for advanced optical networks
- Grids impact education and outreach
  - Providing technologies & resources for training, education, outreach
  - Addressing the Digital Divide
- OSG: a scalable computing infrastructure for science?
  - Strategies needed to cope with increasingly large scale

Grid Project References

- Open Science Grid
  - www.opensciencegrid.org
- Grid3
  - www.ivdgl.org/grid3
- Virtual Data Toolkit
  - www.griphyn.org/vdt
- GriPhyN
  - www.griphyn.org
- iVDGL
  - www.ivdgl.org
- PPDG
  - www.ppdg.net
- CHEPREO
  - www.chepreo.org
- UltraLight
  - ultralight.cacr.caltech.edu
- Globus
  - www.globus.org
- Condor
  - www.cs.wisc.edu/condor
- LCG
  - www.cern.ch/lcg
- EU DataGrid
  - www.eu-datagrid.org
- EGEE
  - www.eu-egee.org

Extra Slides

GriPhyN Goals

- Conduct CS research to achieve vision
  - Virtual Data as unifying principle
  - Planning, execution, performance monitoring
- Disseminate through Virtual Data Toolkit
  - A "concrete" deliverable
- Integrate into GriPhyN science experiments
  - Common Grid tools, services
- Educate, involve, train students in IT research
  - Undergrads, grads, postdocs,
  - Underrepresented groups

iVDGL Goals

- Deploy a Grid laboratory
  - Support research mission of data intensive experiments
  - Provide computing and personnel resources at university sites
  - Provide platform for computer science technology development
  - Prototype and deploy a Grid Operations Center (iGOC)
- Integrate Grid software tools
  - Into computing infrastructures of the experiments
- Support delivery of Grid technologies
  - Hardening of the Virtual Data Toolkit (VDT) and other middleware technologies developed by GriPhyN and other Grid projects
- Education and Outreach
  - Lead and collaborate with Education and Outreach efforts
  - Provide tools and mechanisms for underrepresented groups and remote regions to participate in international science projects

"Virtual Data": Derivation & Provenance

- Most scientific data are not simple "measurements"
  - They are computationally corrected/reconstructed
  - They can be produced by numerical simulation
- Science & eng. projects are more CPU and data intensive
  - Programs are significant community resources (transformations)
  - So are the executions of those programs (derivations)
- Management of dataset dependencies critical!
  - Derivation: Instantiation of a potential data product
  - Provenance: Complete history of any existing data product

- **Previously: Manual methods**
- **GriPhyN: Automated, robust tools**

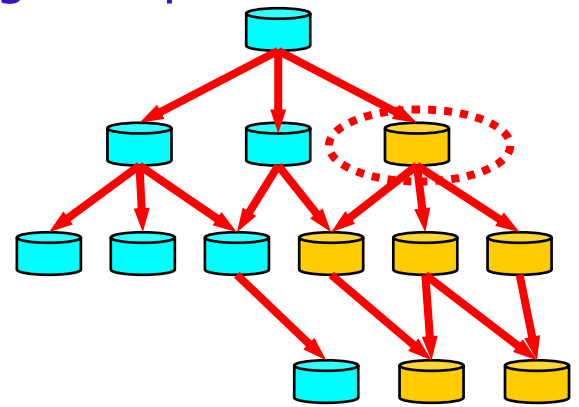

Virtual Data Example: HEP Analysis

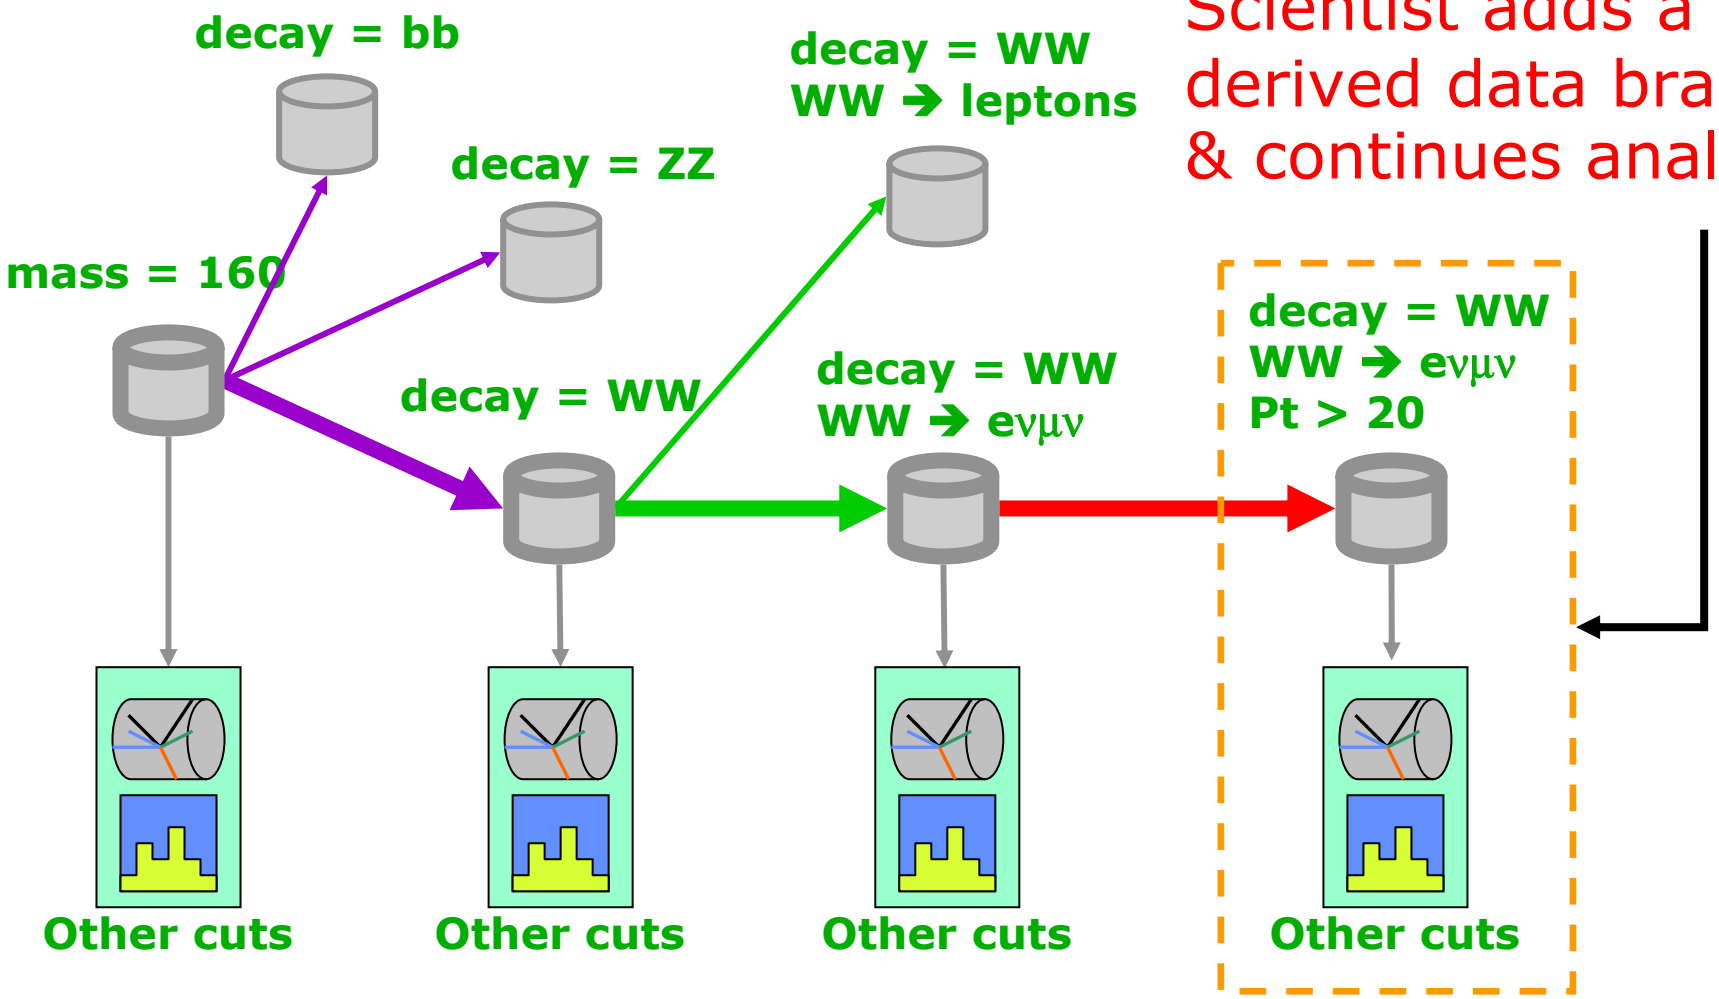

Packaging of Grid Software: Pacman

- Language: define software environments
- Interpreter: create, install, configure, update, verify environments
- Version 3.0.2 released Jan. 2005

- LCG/Scram
- ATLAS/CMT
- CMS DPE/tar/make
- LIGO/tar/make
- OpenSource/tar/make
- Globus/GPT
- NPACI/TeraGrid/tar/make
- D0/UPS-UPD
- Commercial/tar/make

Combine and manage software from arbitrary sources.

"1 button install": Reduce burden on administrators

% pacman –get iVDGL:Grid3

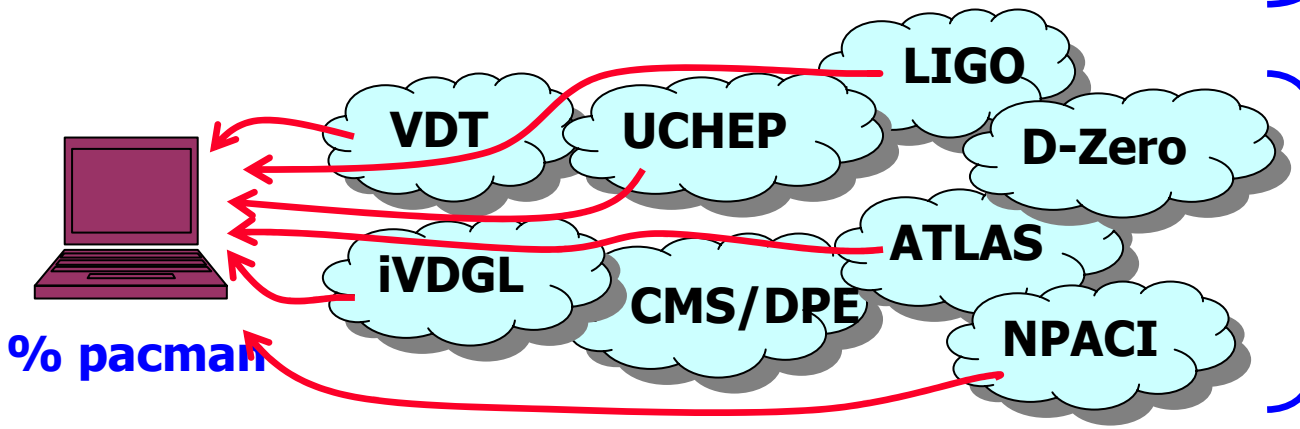

Remote experts define installation/ config/updating for everyone at once

Virtual Data Motivations

"I've found some interesting data, but I need to know exactly what corrections were applied before I can trust it."

"I've detected a muon calibration error and want to know which derived data products need to be recomputed."

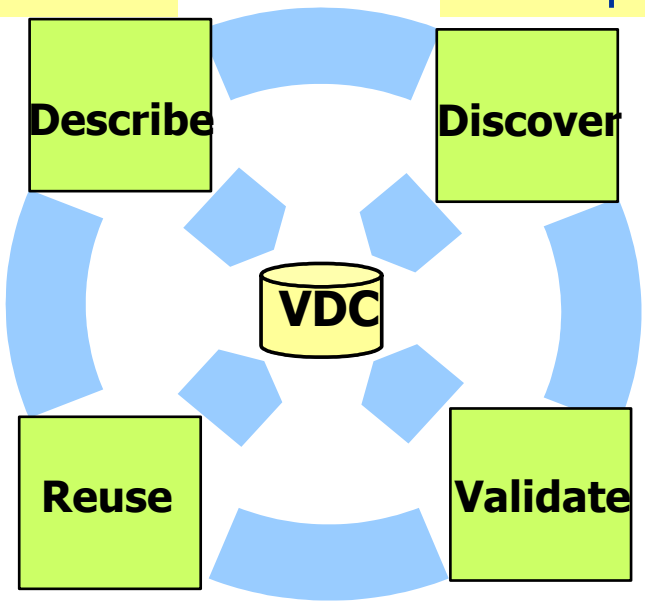

"I want to search a database for 3 muon events. If a program that does this analysis exists, I won't have to write one from scratch."

"I want to apply a forward jet analysis to 100M events. If the results already exist, I'll save weeks of computation."

Background: Data Grid Projects

Driven primarily by HEP applications

- U.S. Funded Projects
  - GriPhyN (NSF)
  - iVDGL (NSF)
  - Particle Physics Data Grid (DOE)
  - UltraLight
  - TeraGrid (NSF)
  - DOE Science Grid (DOE)
  - NEESgrid (NSF)
  - NSF Middleware Initiative (NSF)
- EU, Asia projects
  - EGEE (EU)
  - LCG (CERN)
  - DataGrid
  - EU national Projects
  - DataTAG (EU)
  - CrossGrid (EU)
  - GridLab (EU)
  - Japanese, Korea Projects

- Many projects driven/led by HEP + CS
- Many 10s x $M brought into the field
- Large impact on other sciences, education

"Virtual Data": Derivation & Provenance

- Most scientific data are not simple "measurements"
  - They are computationally corrected/reconstructed
  - They can be produced by numerical simulation
- Science & eng. projects are more CPU and data intensive
  - Programs are significant community resources (transformations)
  - So are the executions of those programs (derivations)
- Management of dataset dependencies critical!
  - Derivation: Instantiation of a potential data product
  - Provenance: Complete history of any existing data product

- **Previously: Manual methods**
- **GriPhyN: Automated, robust tools**

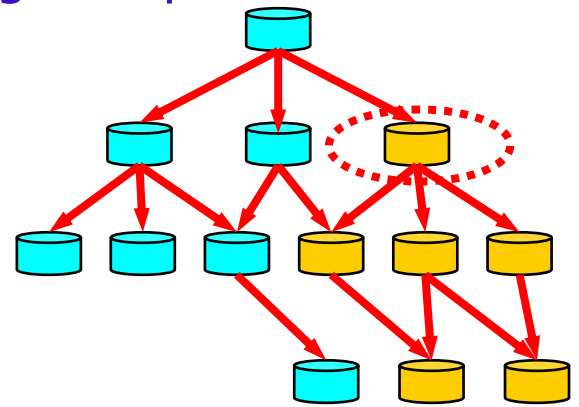

Muon Lifetime Analysis Workflow

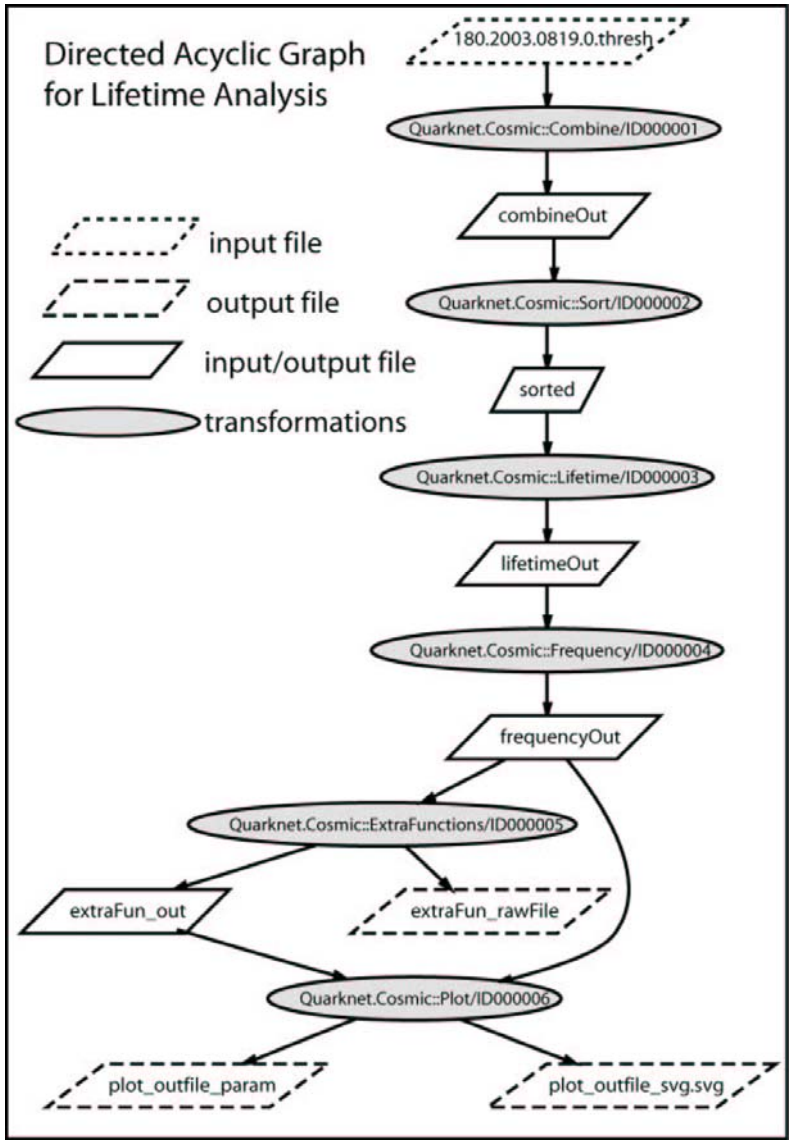

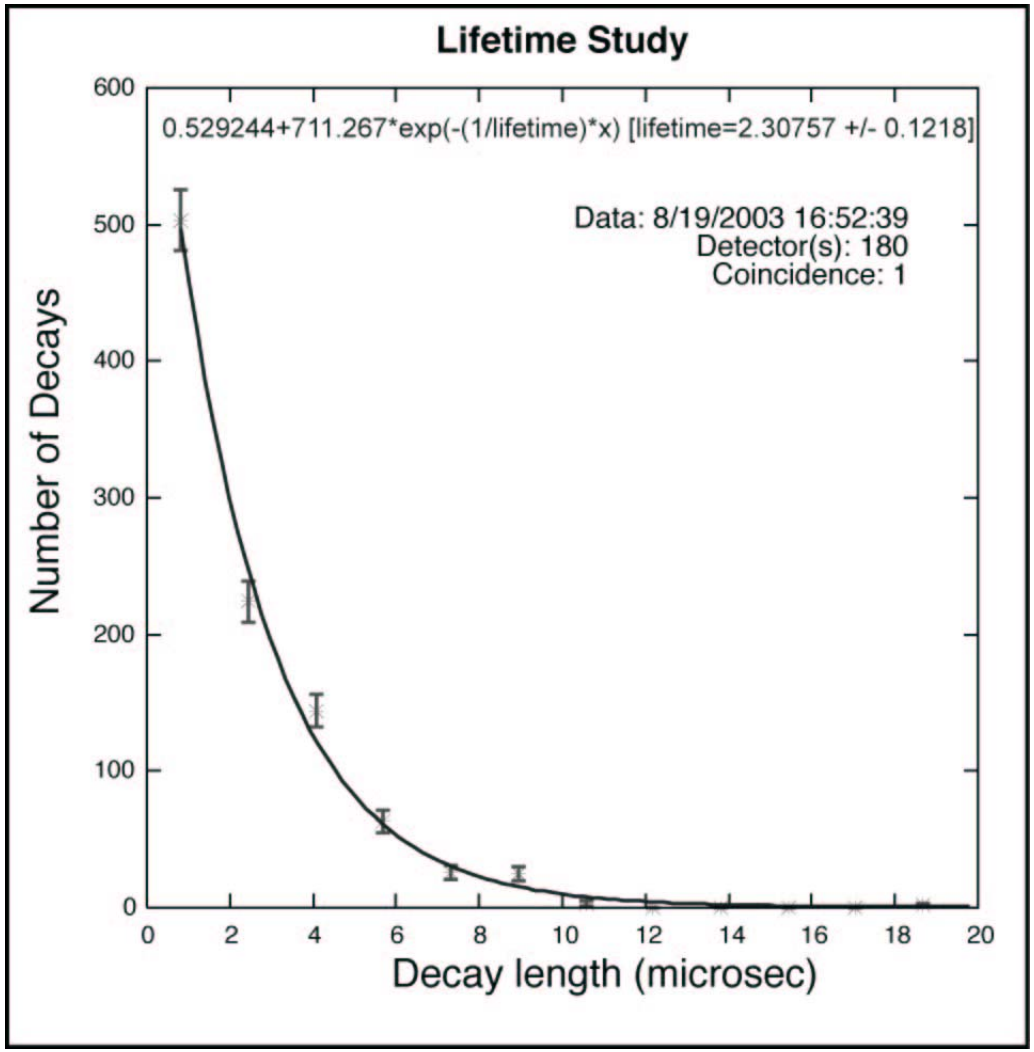

(Early) Virtual Data Language

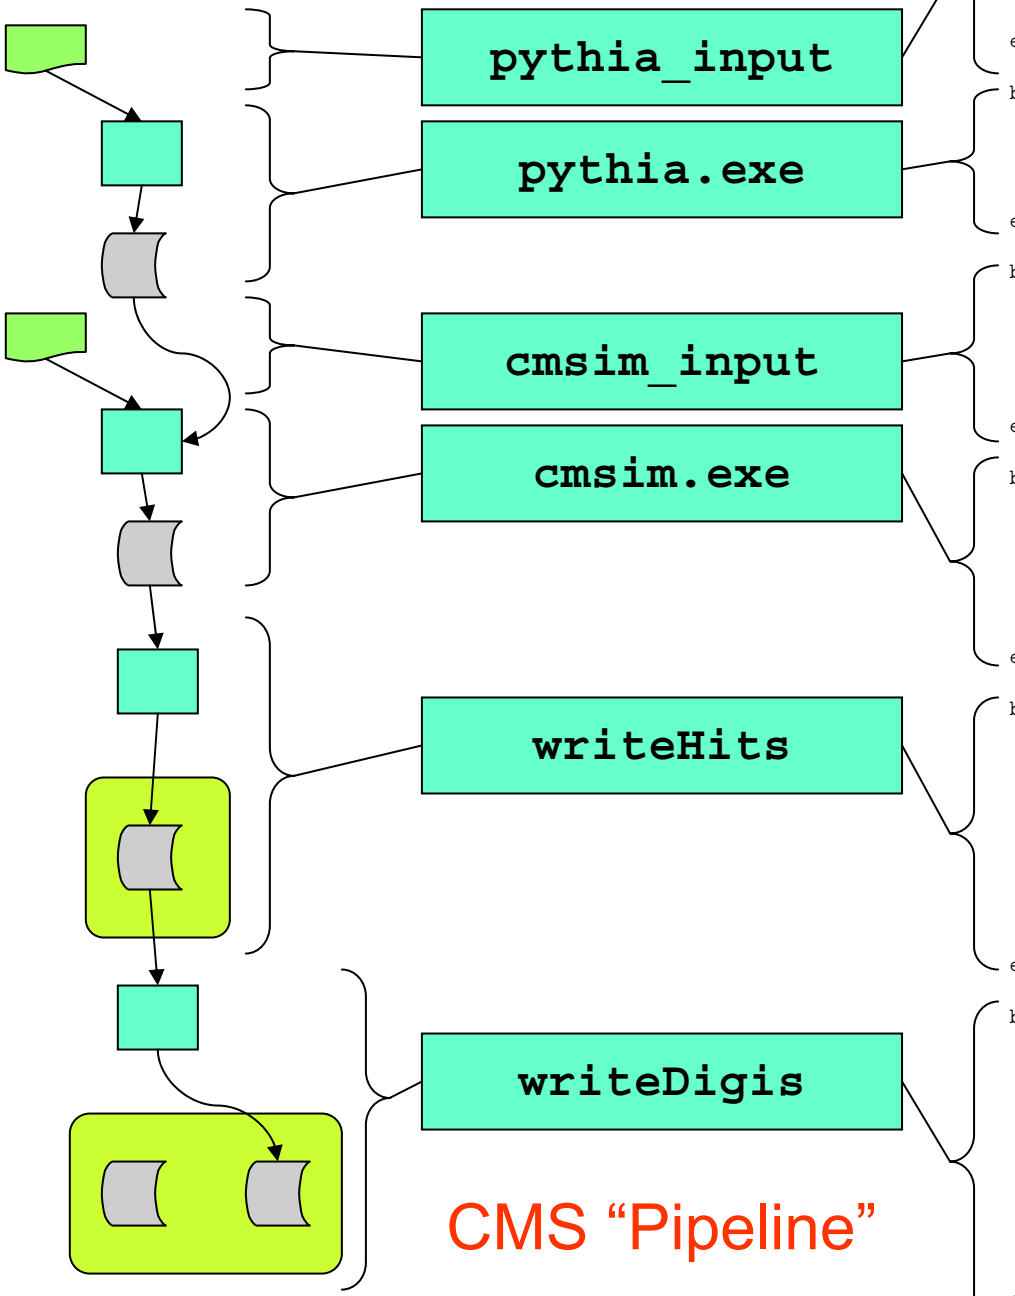

pythia_input

```
begin v /usr/local/demo/scripts/cmkin_input.csh
  file i ntpl_file_path
  file i template_file
  file i num_events
  stdout cmkin_param_file
end
```

pythia.exe

```
begin v /usr/local/demo/binaries/kine_make_ntpl_pyt_cms121.exe
  pre    cms_env_var
  stdin  cmkin_param_file
  stdout cmkin_log
  file o ntpl_file
end
```

cmsim_input

```
begin v /usr/local/demo/scripts/cmsim_input.csh
  file i ntpl_file
  file i fz_file_path
  file i hbook_file_path
  file i num_trigs
  stdout cmsim_param_file
end
```

cmsim.exe

```
begin v /usr/local/demo/binaries/cms121.exe
  condor copy_to_spool=false
  condor getenv=true
  stdin  cmsim_param_file
  stdout cmsim_log
  file o fz_file
  file o hbook_file
end
```

writeHits

```
begin v /usr/local/demo/binaries/writeHits.sh
   condor getenv=true
   pre orca_hits
   file i fz_file
   file i detinput
   file i condor_writeHits_log
   file i oo_fd_boot
   file i datasetname
   stdout writeHits_log
   file o hits_db
end
```

writeDigis

```
begin v /usr/local/demo/binaries/writeDigis.sh
   pre orca_digis
   file i hits_db
   file i oo_fd_boot
   file i carf_input_dataset_name
   file i carf_output_dataset_name
   file i carf_input_owner
   file i carf_output_owner
   file i condor_writeDigis_log
   stdout writeDigis_log
   file o digis_db
end
```

CMS “Pipeline”

QuarkNet Portal Architecture

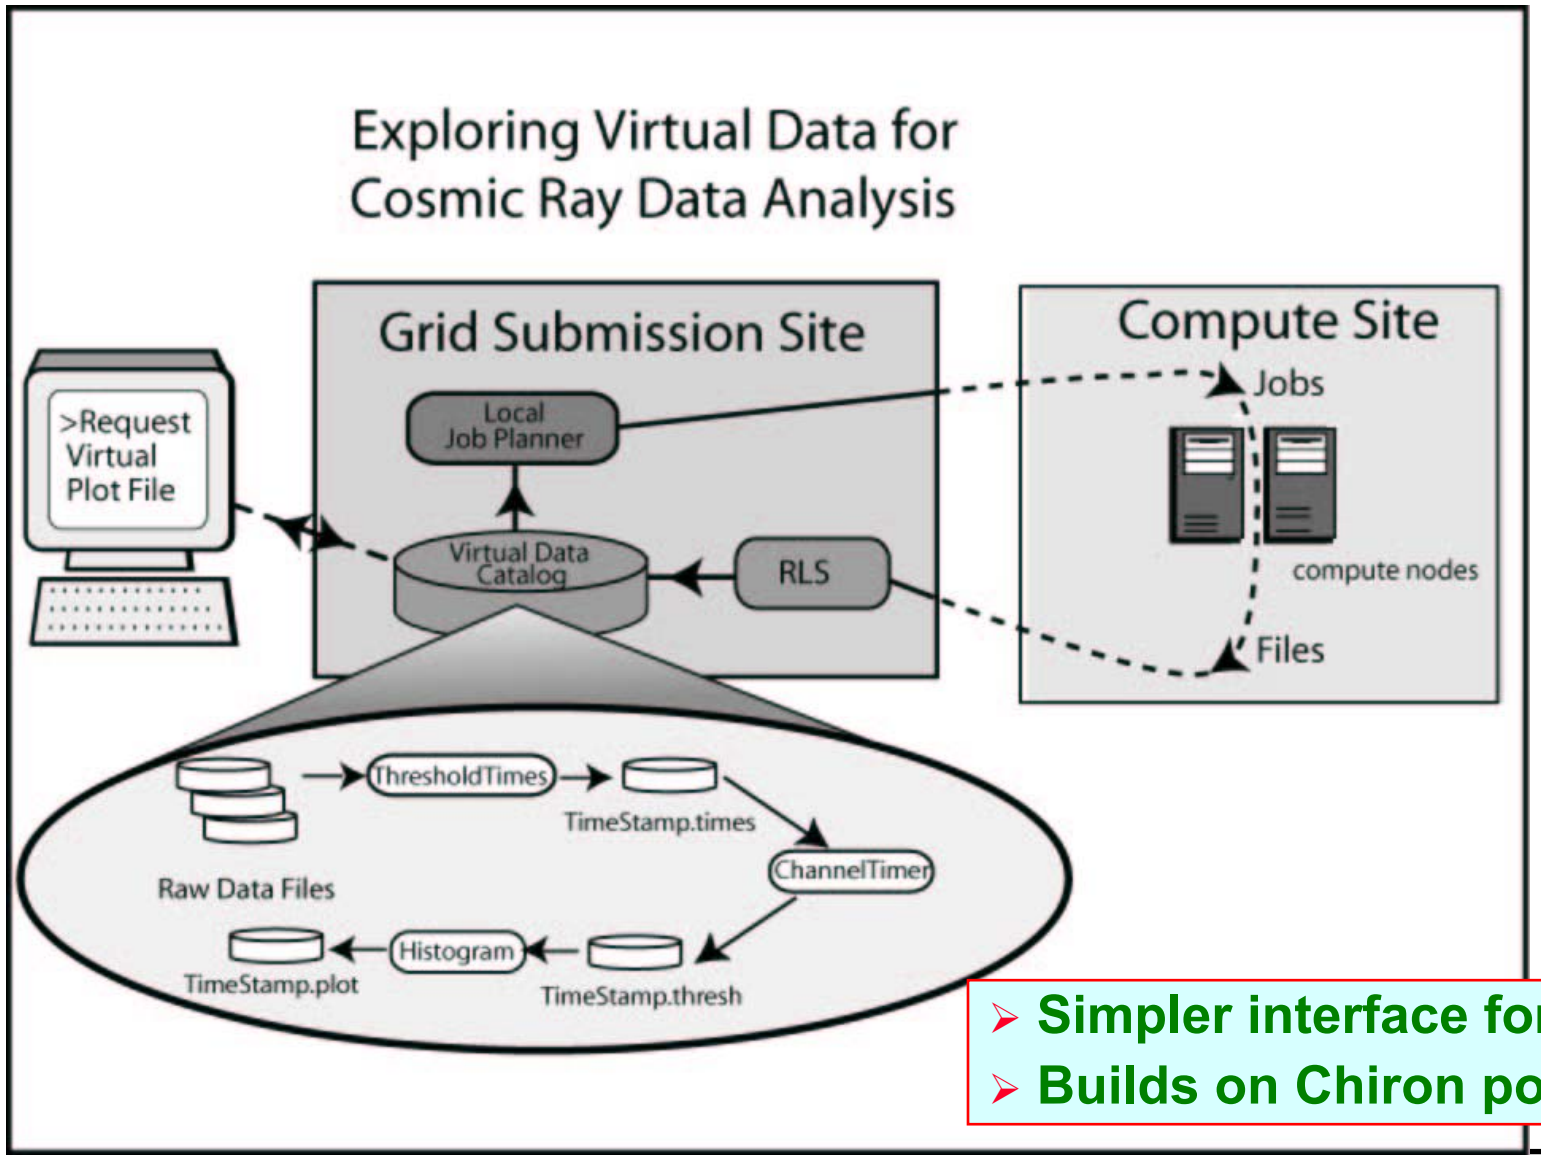

- Simpler interface for non-experts
- Builds on Chiron portal

Integration of GriPhyN and IVDGL

- Both funded by NSF large ITRs, overlapping periods
  - GriPhyN: CS Research, Virtual Data Toolkit (9/2000–9/2005)
  - iVDGL: Grid Laboratory, applications (9/2001–9/2006)
- Basic composition
  - GriPhyN: 12 universities, SDSC, 4 labs (~80 people)
  - iVDGL: 18 institutions, SDSC, 4 labs (~100 people)
  - Expts: CMS, ATLAS, LIGO, SDSS/NVO
- GriPhyN (Grid research) vs iVDGL (Grid deployment)
  - GriPhyN: 2/3 "CS" + 1/3 "physics" (0% H/W)
  - iVDGL: 1/3 "CS" + 2/3 "physics" (20% H/W)
- Many common elements
  - Common Directors, Advisory Committee, linked management
  - Common Virtual Data Toolkit (VDT)
  - Common Grid testbeds
  - Common Outreach effort

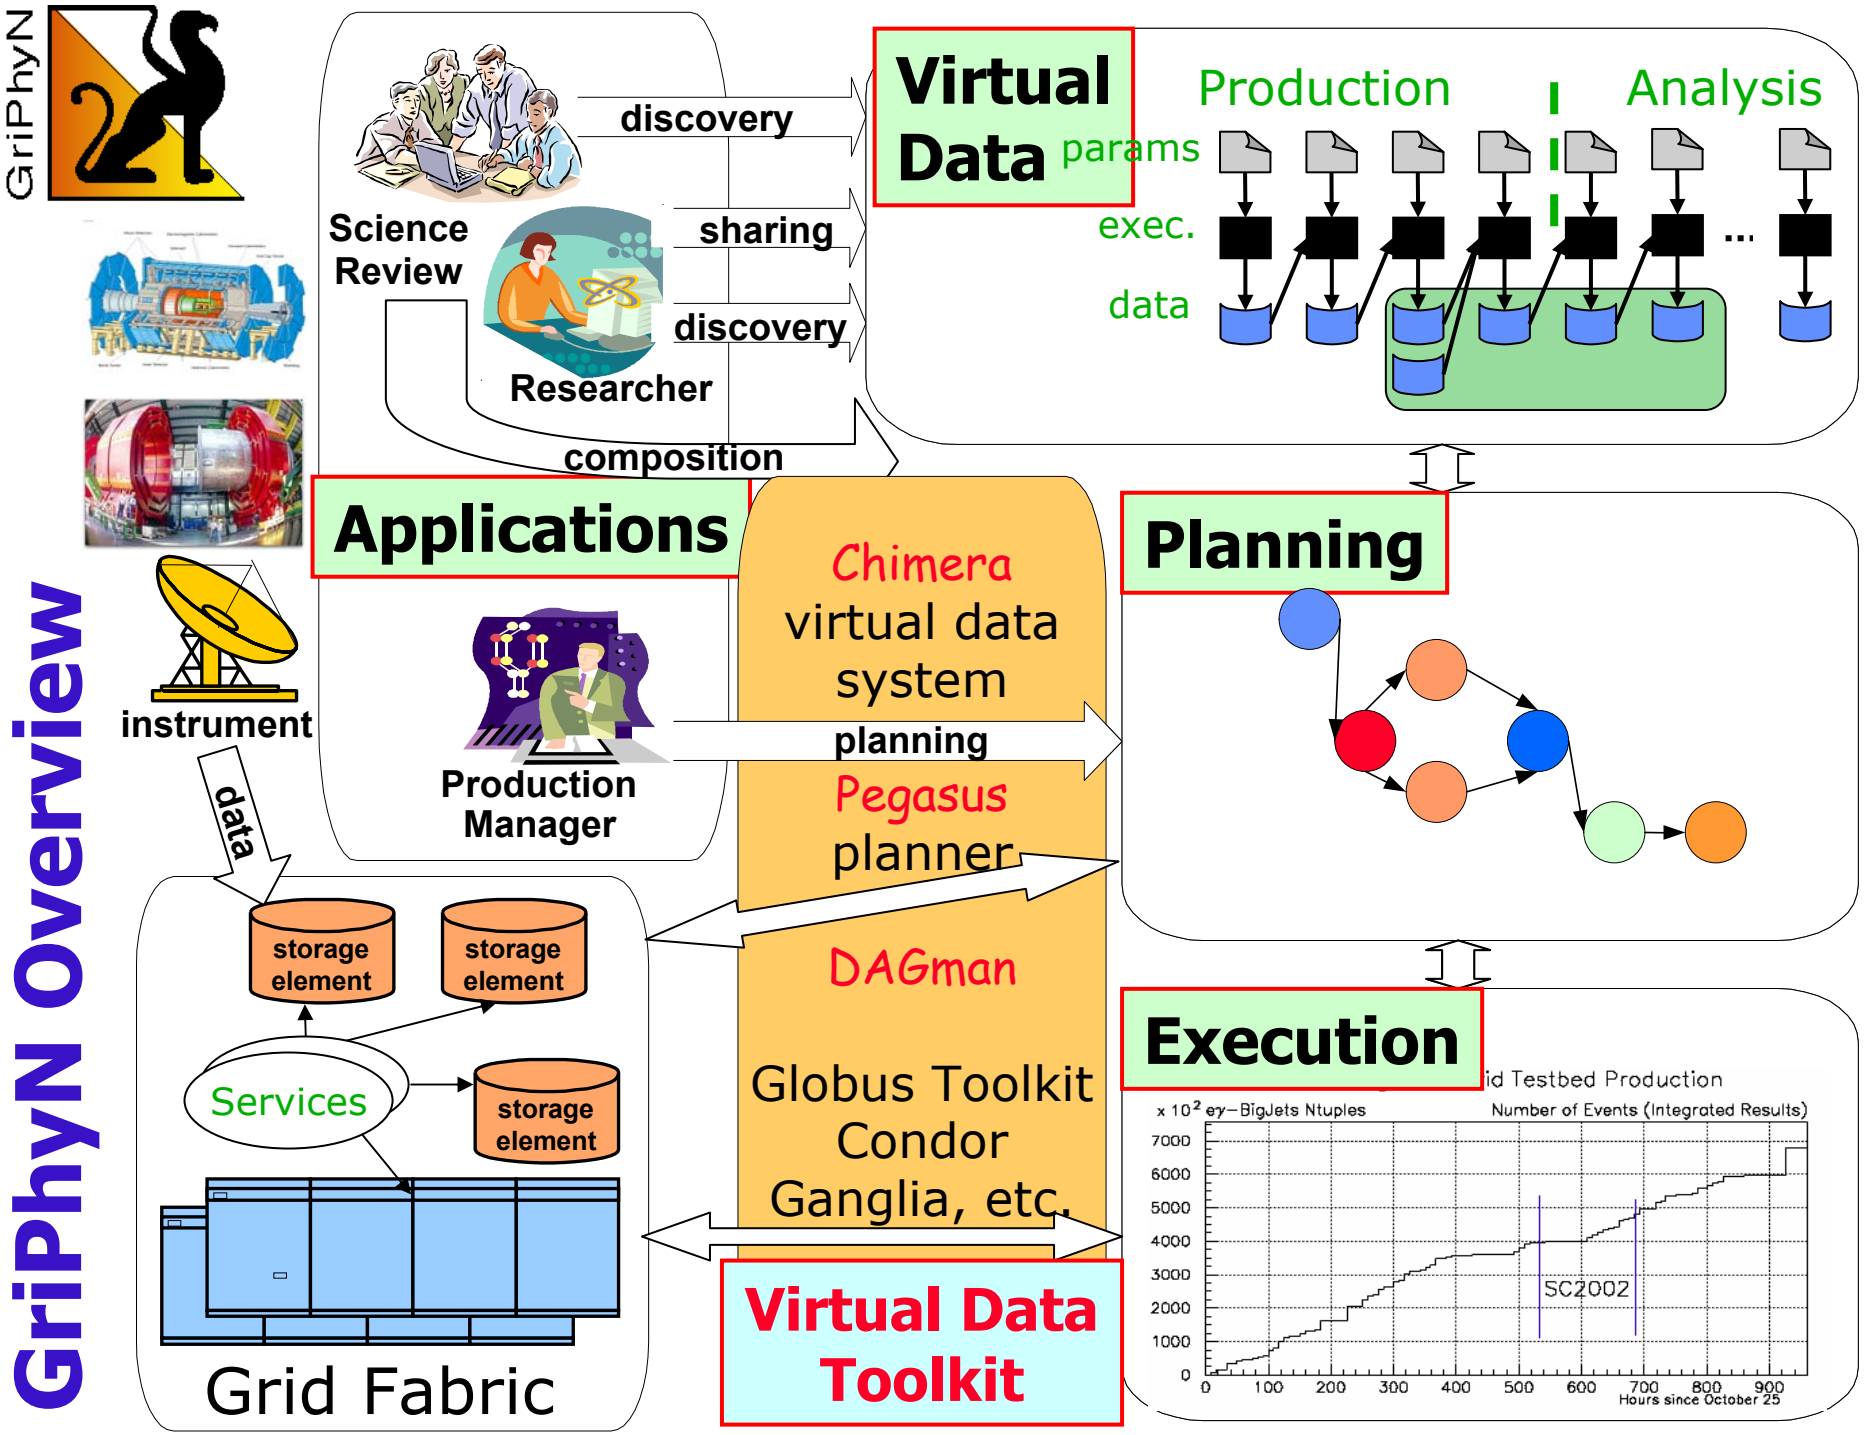
GriPhyN
GriPhyN Overview
Science Review
discovery
sharing
discovery
Researcher
composition
Virtual Data
params
exec.
data
Production
Analysis
Applications
Chimera
virtual data
system
planning
Production
Manager
Planning
instrument
data
Pegasus
planner
DAGman
Globus Toolkit
Condor
Ganglia, etc.
storage element
storage element
Services
storage element
Grid Fabric
Execution
Virtual Data
Toolkit
id Testbed Production
x 10 2 ey−BigJets Ntuples
Number of Events (Integrated Results)
7000
6000
5000
4000
3000
2000
1000
0
0
100
200
300
400
500
600
700
800
900
Hours since October 25
SC2002

Chiron/QuarkNet Architecture

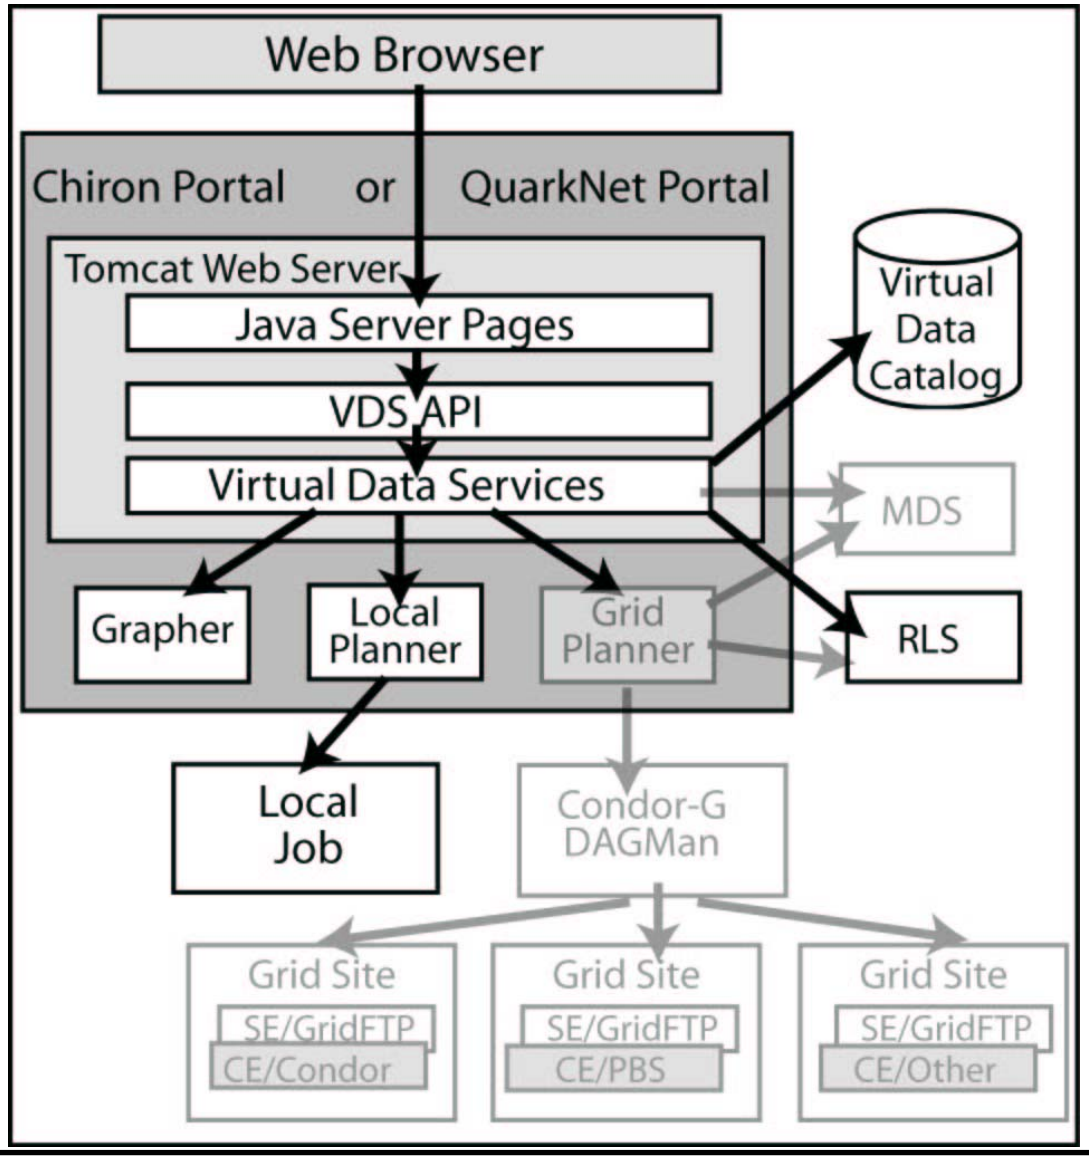

Cyberinfrastructure

"A new age has dawned in scientific & engineering research, pushed by continuing progress in computing, information, and communication technology, & pulled by the expanding complexity, scope, and scale of today's challenges. The capacity of this technology has crossed thresholds that now make possible a comprehensive "cyberinfrastructure" on which to build new types of scientific & engineering knowledge environments & organizations and to pursue research in new ways & with increased efficacy."

[NSF Blue Ribbon Panel report, 2003]

Fulfilling the Promise of Next Generation Science

Our multidisciplinary partnership of physicists, computer scientists, engineers, networking specialists and education experts, from universities and laboratories, has achieved tremendous success in creating and maintaining general purpose cyberinfrastructure supporting leading-edge science.

But these achievements have occurred in the context of overlapping short-term projects. How can we ensure the survival of valuable *existing* cyber-infrastructure while continuing to address *new* challenges posed by frontier scientific and engineering endeavors?

Production Simulations on Grid3

US-CMS Monte Carlo Simulation

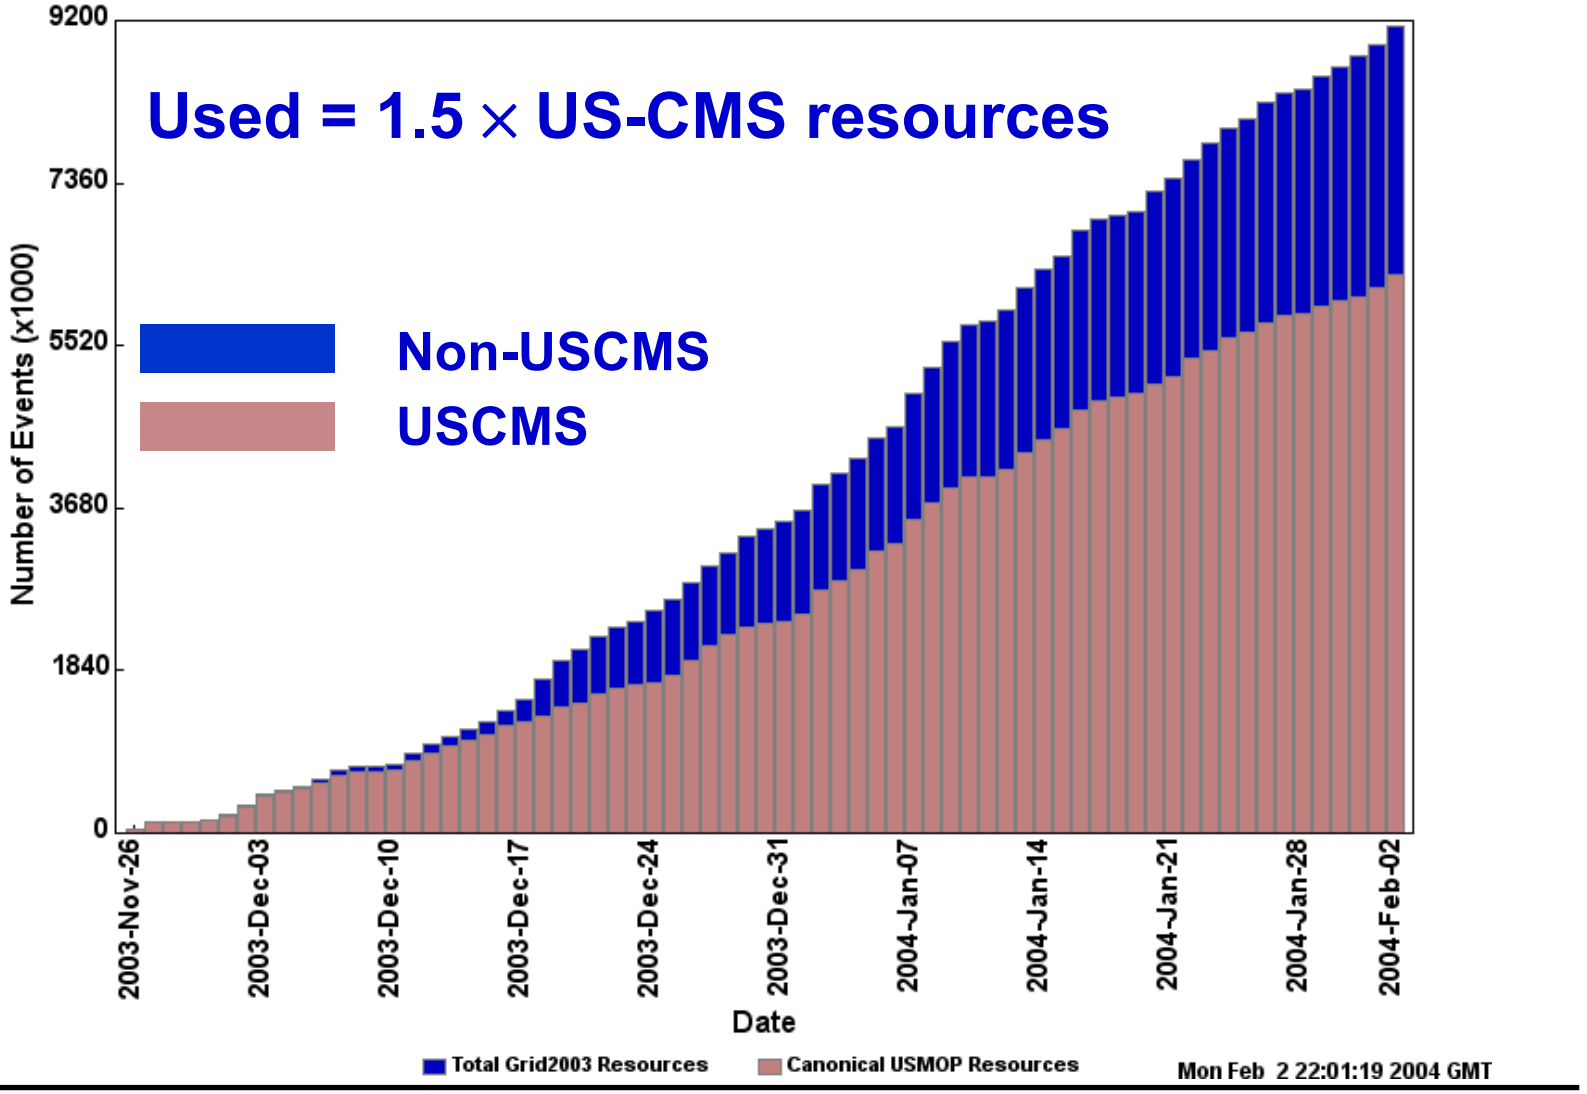

Components of VDT 1.3.5

- Globus 3.2.1
- Condor 6.7.6
- RLS 3.0
- ClassAds 0.9.7
- Replica 2.2.4
- DOE/EDG CA certs
- ftsh 2.0.5
- EDG mkgridmap
- EDG CRL Update
- GLUE Schema 1.0
- VDS 1.3.5b
- Java
- Netlogger 3.2.4
- Gatekeeper-Authz
- MyProxy1.11
- KX509
- System Profiler
- GSI OpenSSH 3.4
- Monalisa 1.2.32
- PyGlobus 1.0.6
- MySQL
- UberFTP 1.11
- DRM 1.2.6a
- VOMS 1.4.0
- VOMS Admin 0.7.5
- Tomcat
- PRIMA 0.2
- Certificate Scripts
- Apache
- jClarens 0.5.3
- New GridFTP Server
- GUMS 1.0.1

Collaborative Relationships: A CS + VDT Perspective