The Particle Physics Data Grid Collaboratory Pilot

www.ppdg.net

Richard P. Mount
SLAC

March 19, 2003

The HENP Collaboratory Past

- **30 years ago an HEP "collaboratory" involved:**
  - Air freight of bubble chamber film (*e.g.* CERN to Cambridge)
- **20 years ago:**
  - Tens of thousands of tapes
  - 100 physicists from all over Europe (or US)
  - Air freight of tapes, 300 baud modems
- **10 years ago:**
  - Tens of thousands of tapes
  - 500 physicists from US, Europe, USSR, PRC …
  - 64k bps leased lines and air freight

www.ppdg.net

The Collaboratory Present and Future

- **Present:**
  - Tens of thousands of tapes
  - 500 physicists from US, Europe, Japan, FSU, PRC …
  - Dedicated intercontinental links at up to 155/622 Mbps
  - Home brewed, experiment-specific, data/job distribution software (if you're lucky)
  - Plus some useful Grid middleware components
- **Future (~2010):**
  - Tens of thousands of tapes
  - 2000 physicists from, worldwide collaboration
  - Many links at 2.5/10 Gbps
  - The Grid must exist (but what will it be?)

www.ppdg.net

Grid Fabric and Motivation

BaBar Example

www.ppdg.net

BaBar Computing Outlook

	A	B	C	D	E	F	G	H	I
1		Unit	1999	2000	2001	2002	2003	2004	2005
2									
3	**Estimates (Nov. 2001)**								
4	**Peak Luminosity**	**10^{33} cm^{-2} sec^{-1}**	**1**	**2**	**4.4**	**5**	**7.5**	**10**	**13**
5	**Yearly Integrated Luminosity**	**fb^{-1}**	**3**	**23**	**41**	**39**	**62.6**	**66.1**	**120.1**
6	**Contingency factor**		**1**	**1**	**1**	**1.2**	**1.2**	**1.2**	**1.2**
7	Peak Luminosity with contingenc	10^{33} cm^{-2} sec^{-1}	1	2	4.4	6	9	12	15.6
8	Yearly Integrated Luminosity with	fb^{-1}	3	23	41	46.8	75.12	79.32	144.12
9									

LumChart | **Luminosity** | Analysis | Production | Unit Costs | Effective Subsystem Costs | Costs

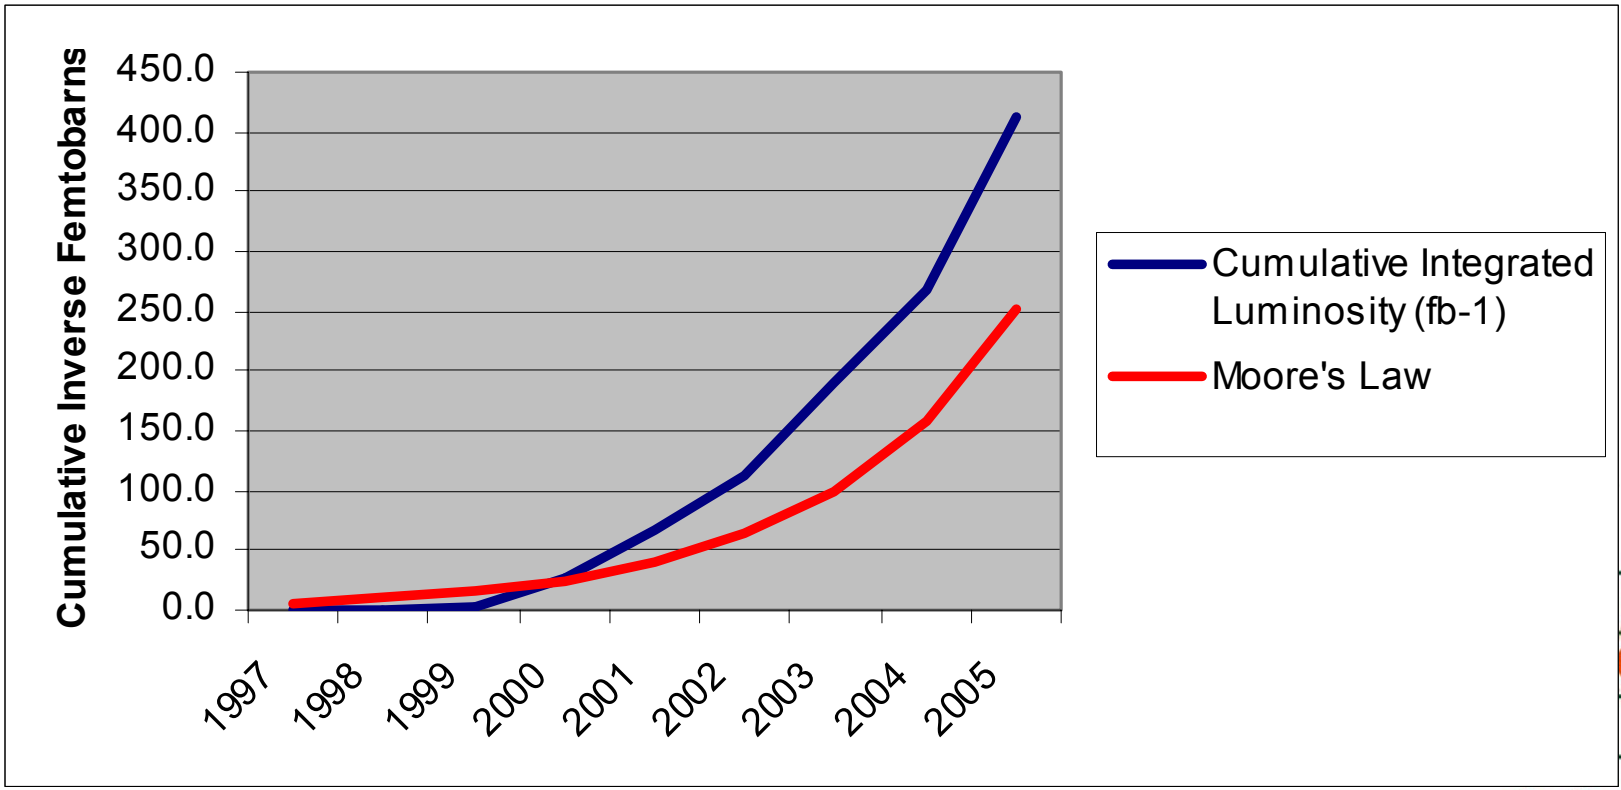

dg.net

P P D G

BaBar Offline Computing Equipment
Bottom-up Cost Estimate (January 2002)

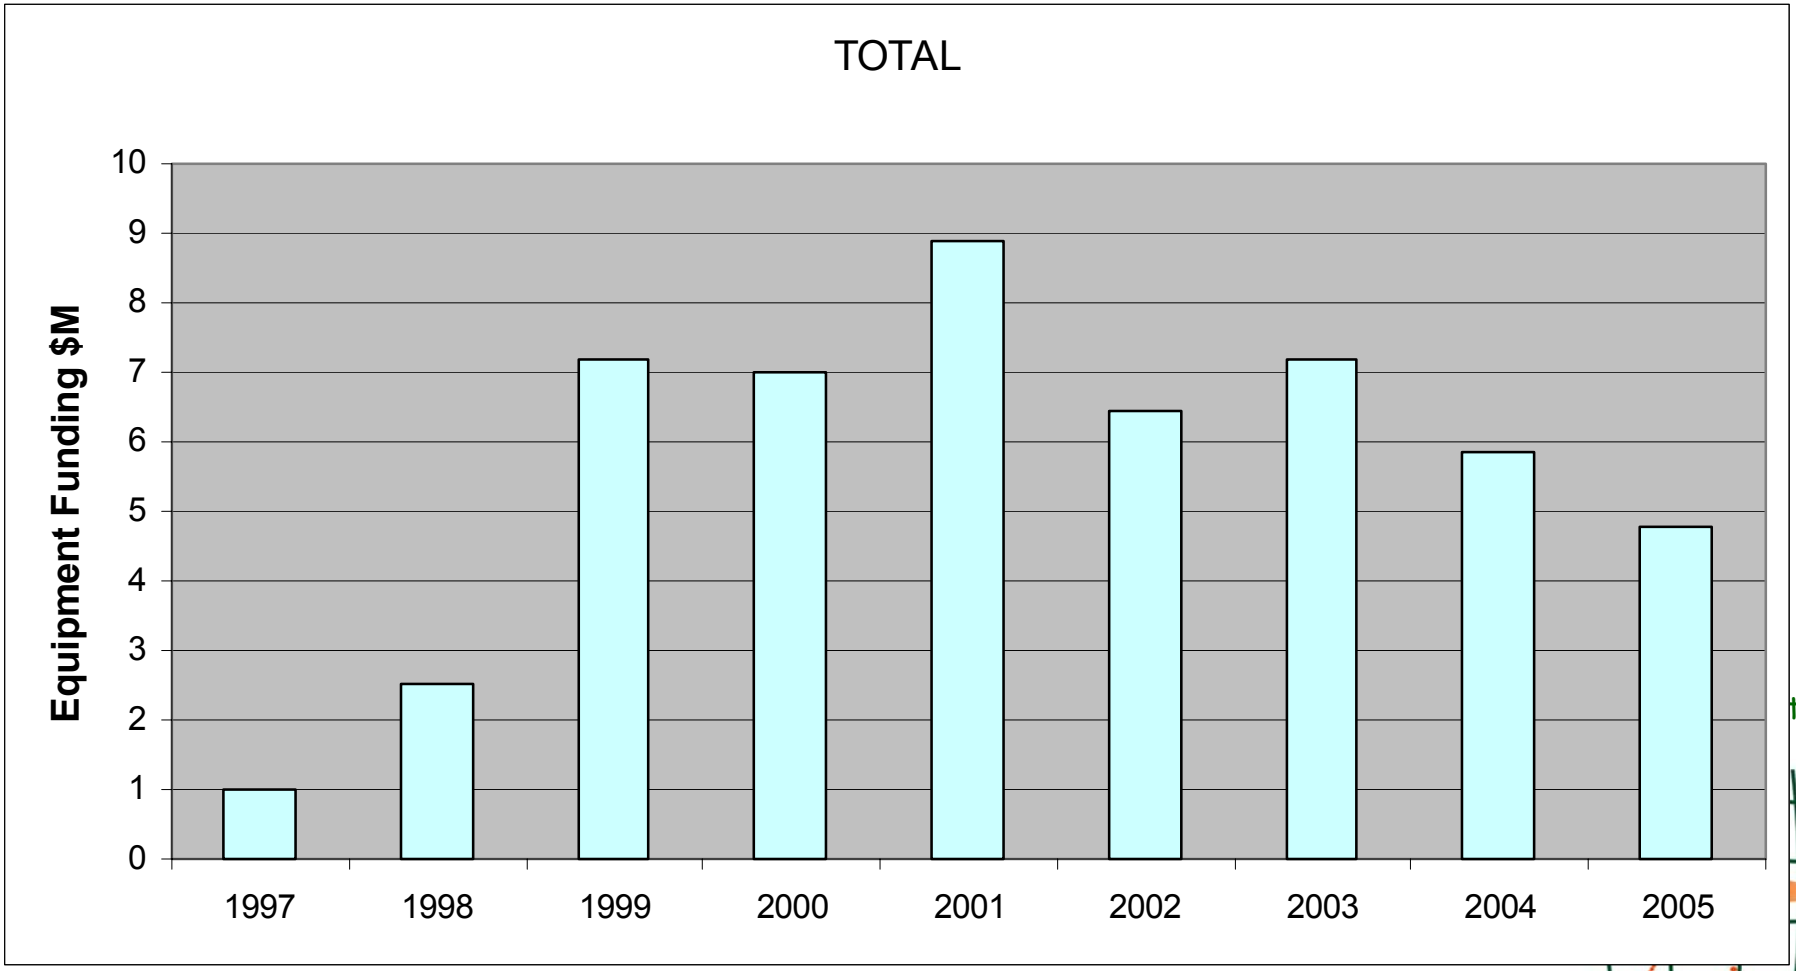

PPDG

BaBar Computing Resources

The "Tier-A" Complex:

- An MOU-defined resource for the whole collaboration
- Optimized for heavy, data-intensive use
- SLAC (US), CCIN2P3 (France), RAL (UK), INFN-Padova (Italy), [GridKa (Germany)]
- Over 300 TB dedicated disk
- Over 3 Tflops (over 3000 processors)
- Robotic mass-storage systems
- Annual membership: ≥ $500k hardware investment for BaBar per site

Simulation Complex

- A resource for the whole collaboration
- Optimized for compute-intensive simulation
- Universities and labs in UK, France, Canada, US

Tier C sites

- Universities with significant local resources, normally serving local users

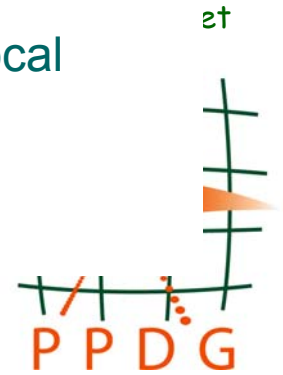

BaBar Tier-A Centers

A component of the Fall 2000 BaBar Computing Model)

- **Offer resources at the disposal of BaBar;**
- **Each provides tens of percent of total BaBar computing/analysis need;**
  - 50% of BaBar computing investment was in Europe in 2002
- **CCIN2P3, Lyon, France in operation for 2 years;**
- **RAL, UK in operation for 1 year**
- **INFN-Padova, Italy in operation for 0.6 years;**
- **GridKA, Karlsruhe, Germany coming on line – will only offer Grid-based services.**

MoU-Defined Grid

Draft, January 12, 2001

BaBar TierA Computing Center MoU

Addendum for a UK Tier-A centre

1. The computing center at RAL will become a BABAR Tier A center.
2. The initial period of this agreement is for four years, starting in 2002. Termination following this initial period may be effected by either party upon two years written notice. In the event of termination, the UK will not invest any more in BABAR computing at RAL, but the data access will be guaranteed for one additional year.
3. Schedule
   The goal is for complete Tier A operation at RAL in early 2002.

Outline of Planned Facilities

net

Year	Funding (£K)	Total CPUs (30 SpecInt 95)	Total Disk (usable TB)

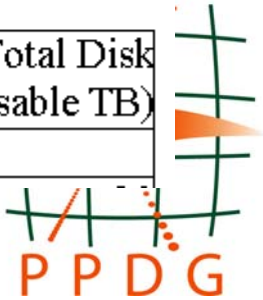

PPDG

www.ppdg.net

Particle Physics Data Grid Collaboratory Pilot

Computer Science Groups: Condor, Globus, SRM, SRB

Nuclear Physics Experiments: STAR, TJNAF

Lattice QCD: TJNF

High Energy Physics: ATLAS, BaBar, CMS, D0

PI's: Richard Mount, SLAC, Miron Livny, Wisconsin, Harvey Newman, Caltech

Steering Committee: John Huth, Harvard (ATLAS), Tim Adye, RAL (BaBar), Lothar Bauerdick, FNAL (CMS), Lee Lueking FNAL (D0), Chip Watson, TJNAF, Jerome Lauret, BNL (STAR), Miron Livny, Wisconsin (Condor), Jennifer Schopf, ANL (Globus), Ian Foster, ANL (Globus), Reagan Moore, SDSC (SRB), Arie Shoshani, LBNL (SRM)
Coordinators: Ruth Pordes, FNAL, Doug Olson, LBNL
Liaisons: Paul Avery (iVDGL), Larry Price (HICB), Mike Wilde(GriPhyN), Torre Wenaus, Ian Bird (LCG)

Program Managers: Mary Anne Scott (MICS), Irwin Gaines (HENP), Steve Steadman (NP)

Experiment data handling requirements:

- Petabytes of storage, Teraops of computing, Thousands of users,
- Hundreds of institutions, 10+ years of analysis ahead

Focus of PPDG:

- Vertical Integration of Grid technologies into Applications ongoing work
- Deployment, hardening and extensions of common Grid services and standards – data replication, storage and job management, monitoring and planning.
- Interdisciplinary teams of physicists, engineers and computer scientists
- Using and extending grid technologies ranging across architecture, integration, deployment and robustness
- Driven by demanding end-to-end applications of experimental physics
- Enabling new scales of research in experimental physics and experimental computer science

www.ppdg.net

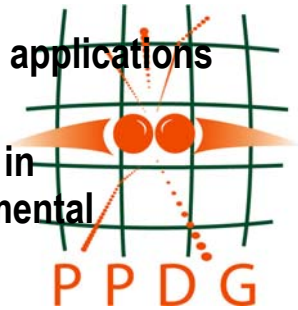

The Novel Ideas

- End to end integration and deployment of experiment applications using existing and emerging Grid services.
- Deployment of Grid technologies and services in production (24x7) environments with stressful performance needs.
- Collaborative development of Grid middleware and extensions between application and middleware groups – leading to pragmatic and least risk solutions.
- HENP experiments extend their adoption of common infrastructures to higher layers of their data analysis and processing applications.
- Much attention paid to integration, coordination, interoperability and interworking with emphasis on incremental deployment of increasingly functional working systems.

www.ppdg.net

PPDG Accomplishments

Application "Grids" with incremental capabilities:

- **Data Replication for BaBar** - Terabyte stores replicated from California to France and England.
- **Replication and Storage Management for STAR** and **JLAB** - development and deployment of standard API and interoperable implementations.
- Production Simulation Grids for **ATLAS** and **CMS, STAR** distributed analysis jobs.
- **Data Transfer, Job and Information Management for D0** - GridFTP integrated with SAM; Condor-G job scheduler, MDS resource discovery all integrated with SAM.

Initial Security Infrastructure for Virtual Organizations:

- **PKI certificate management**, policies and trust relationships (using DOE Science Grid and Globus)
- Authorization mechanisms – standard callouts for Local Center Authorization for Globus, EDG,
- Prototyping secure credential stores
- Engagement of site security teams.

Data and Storage Management:

- Robust data transfer over heterogeneous networks using standard protocols: GridFTP, bbcp
- Distributed Data Replica management: SRB, SAM, SRM
- Common Storage Management interface and services across diverse implementations: SRM - HPSS, Jasmine, Enstore

Job Planning, Execution and Monitoring:

- Job scheduling based on resource discovery and status - condor-g and extensions
- Retry and Fault Tolerance in response to error conditions - hardened gram, gass-cache, ftsh,
- Distributed monitoring infrastructure for resource discovery, resource and job information - MDS, Monalisa, Hawkeye

www.ppdg.net

Prototypes and Evaluations

- Grid enabled physics analysis tools
- End to end troubleshooting and fault handling
- Cooperative Monitoring of Grid, Fabric, Applications

PPDG Collaboratory accomplishment timeline

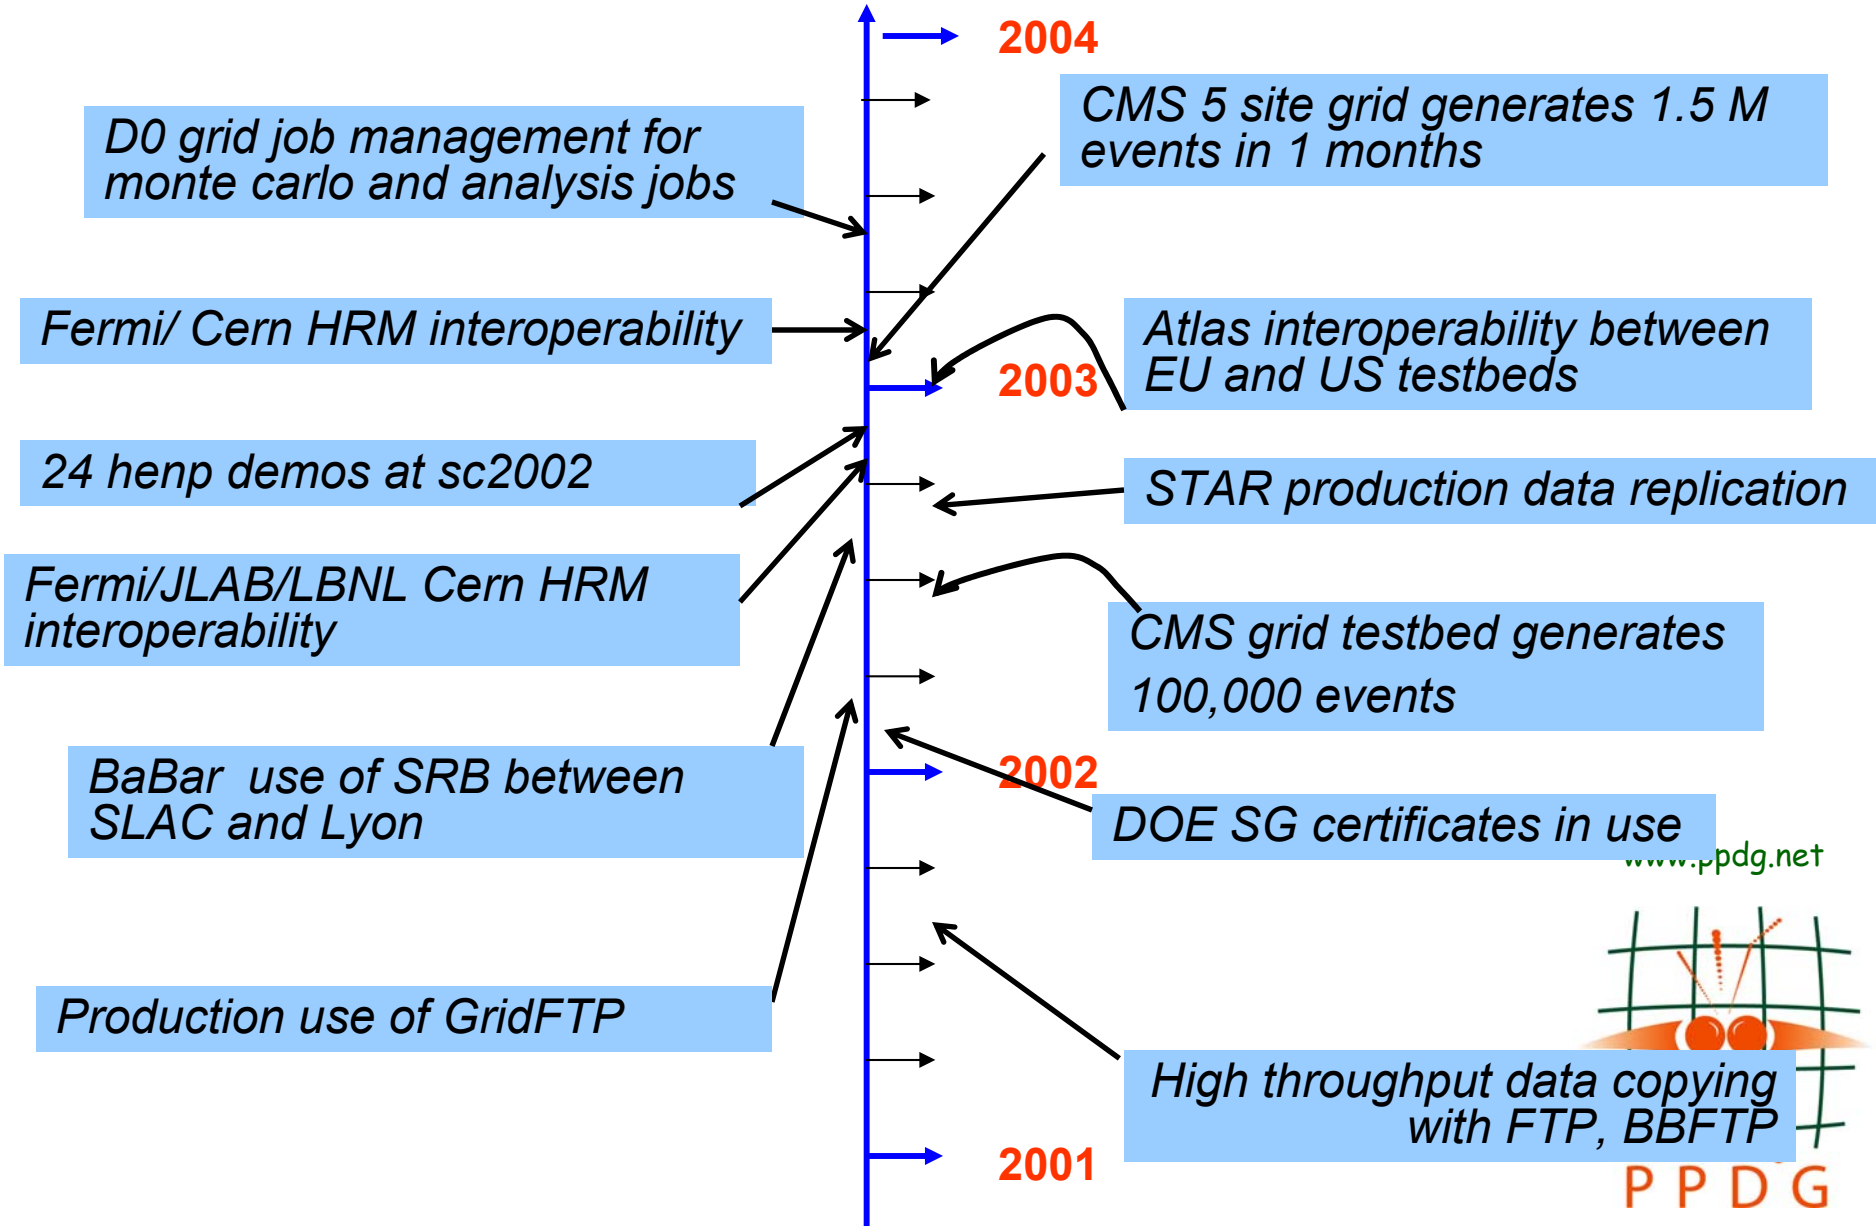

Technologies, Tools and Applications

Who uses what?

	Atlas	BaBar	CMS	D0	STAR	TJNAF	TJNAF-LQCD
Globus GridFTP	x	bbcp	x	x	x	x	
Globus MDS	x		x	x			
Globus GSI	x	x	x (+KCA)	x (+KCA)	x	x	Java
Condor-G/GRAM	x/VDT	x/EDG	x/VDT	x			
SRM			x	x	x	x	
SRB		x	x				
Distributed Data Management	MAGDA	Bbserver++	MOP	SAM	x	x	x
Analysis prototypes		JAS	Caigee Clarens	root/sam			
Automated Data Replication	x	x		x	x	x	x
Applications Grid	p	p/EDG	x	x/sam			
DOE SG Certs	x	x	x	x	x	x	x

www.ppdg.net

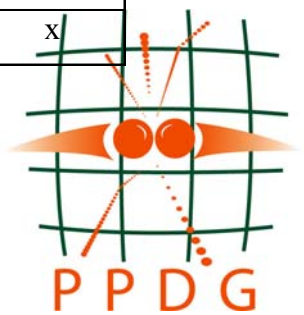

The Future of the Grid

Increasing underlying complexity

- More components
- More services, especially higher level services (e.g. to physics analysis)

Robustness will be essential

- Robustness is not just a bolt-on addition to the core computer science, it *IS* a core computer science

Troubleshooting tools will be essential

- Likely to have major influence on Grid middleware design and architecture

www.ppdg.net

PPDG Collaborations and Interactions

US Physics Grid Projects

GriPhyN, iVDGL, PPDG – the Trillium + Virtual Data Toolkit

EU-US Physics Grids

LHC Computing Grid – ATLAS, CMS

European Data Grid - BaBar, D0, ATLAS, CMS

High Energy Physics Intergrid Coordination Board

SciDAC projects	**Interaction**
Earth System Grid II	SRM
DOE Science Grid	CA, RA
A High-Performance Data Grid Toolkit	Globus Toolkit
Storage Resource Management for Data Grid Applications	SRM
Security and Policy for Group Collaboration	GSI, CAS
Scientific Data Management Center	STAR SDM
Bandwidth Estimation: Measurement Methodologies and Applications	IEPM-BW
A National Computational Infrastructure for Lattice Gauge Theory	TJNAF/LQCD
Distributed Monitoring Framework	Netlogger, Glue Schema

www.ppdg.net

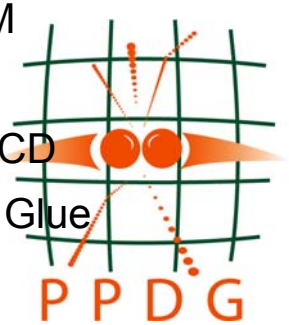

Challenge and Opportunity

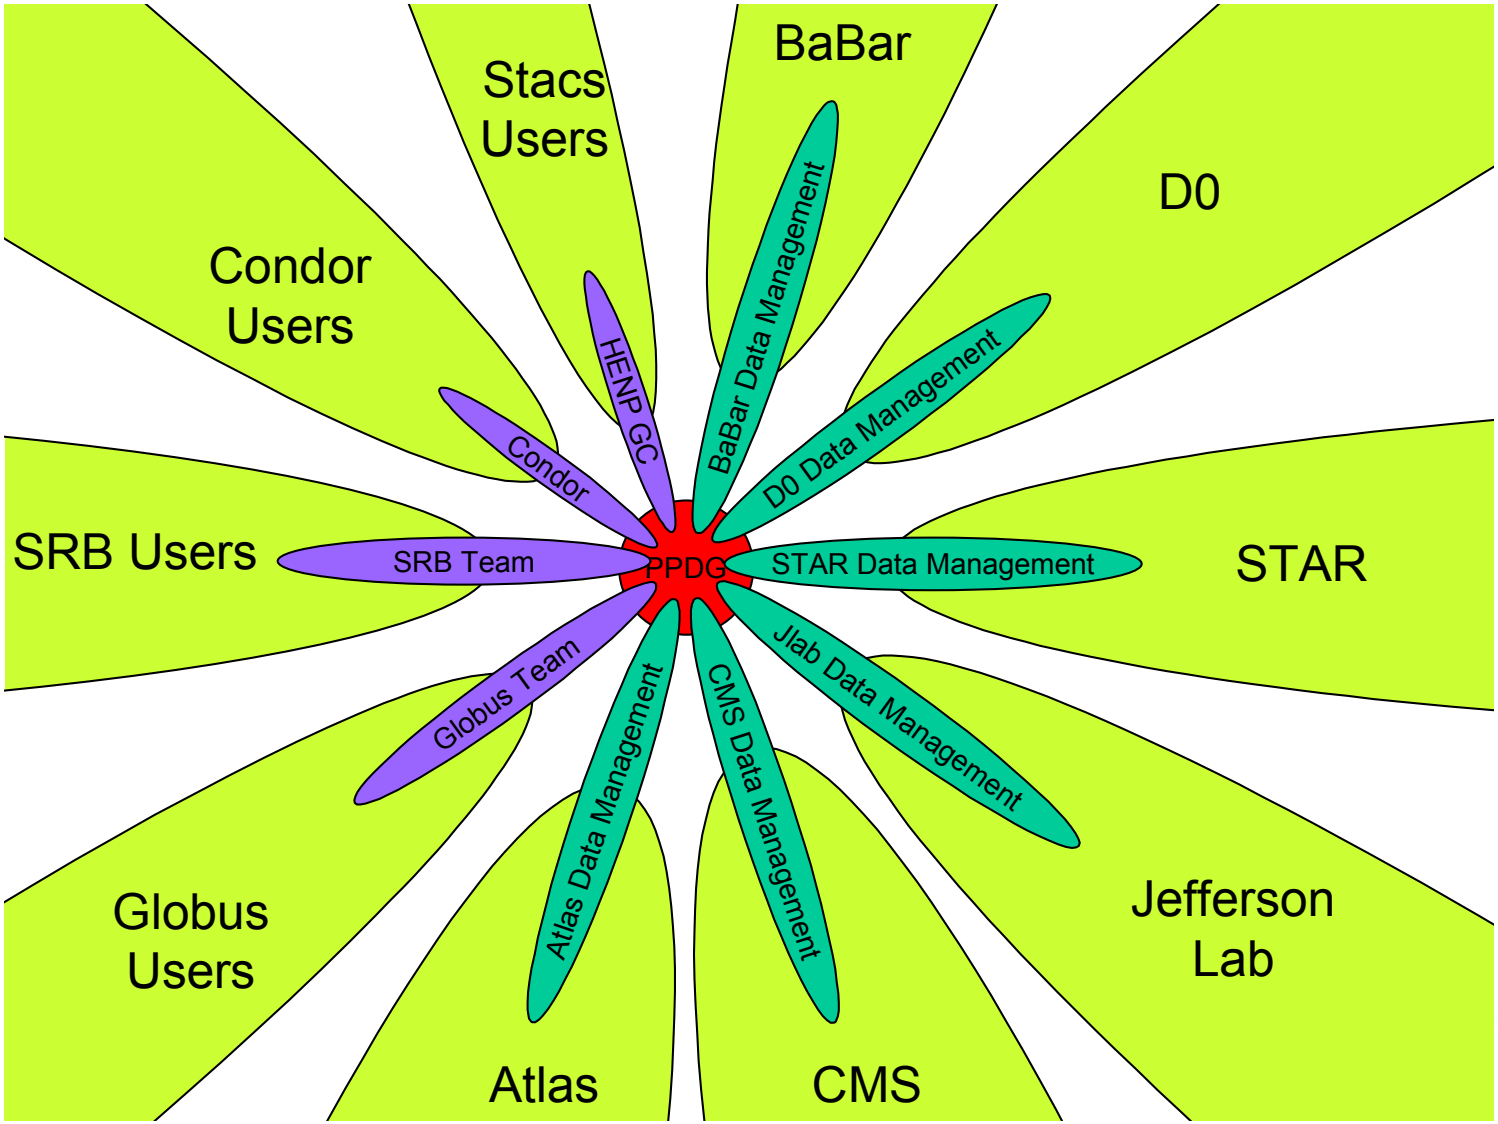

www.ppdg.net

PPDG

PPDG World

An Experiment

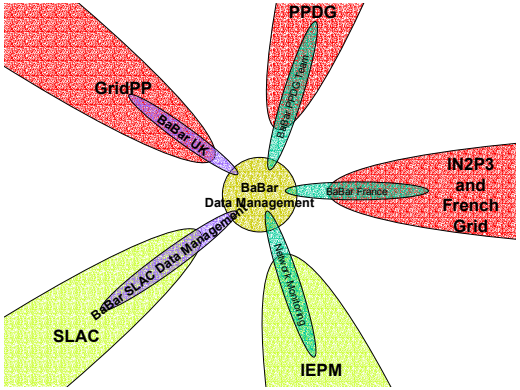

HENP Grid

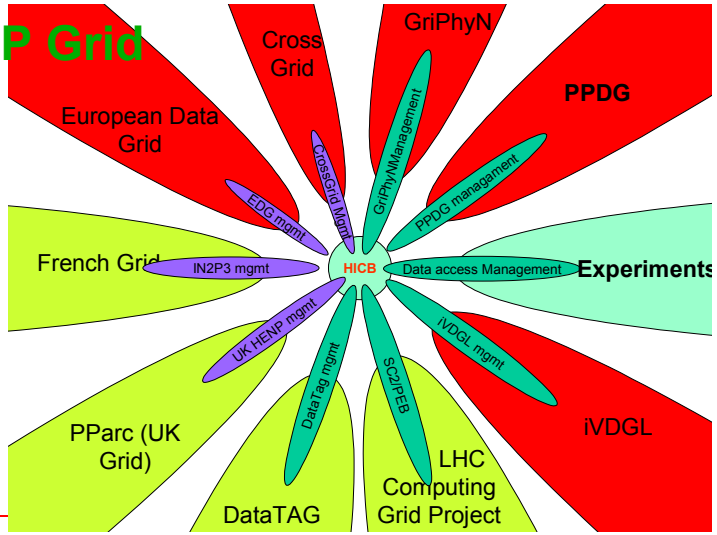

SciDAC connections

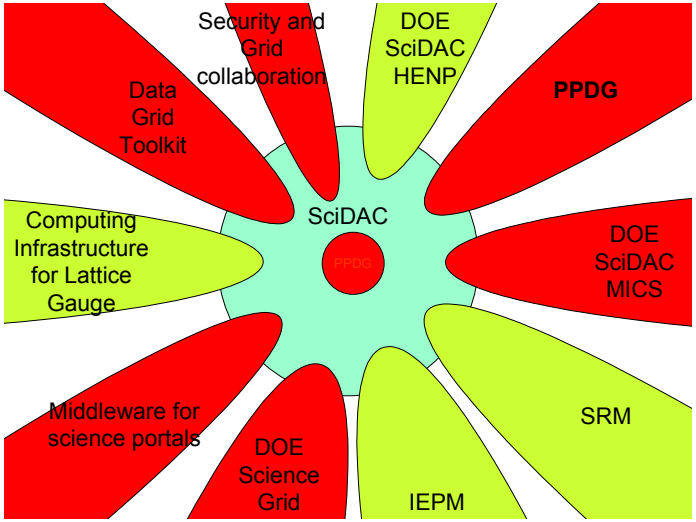

PPDG

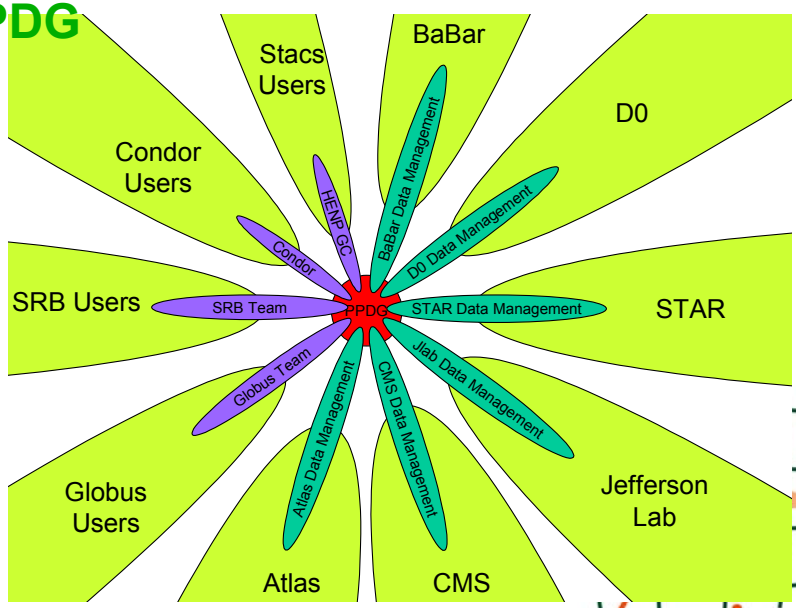

HENP Grid Timelines

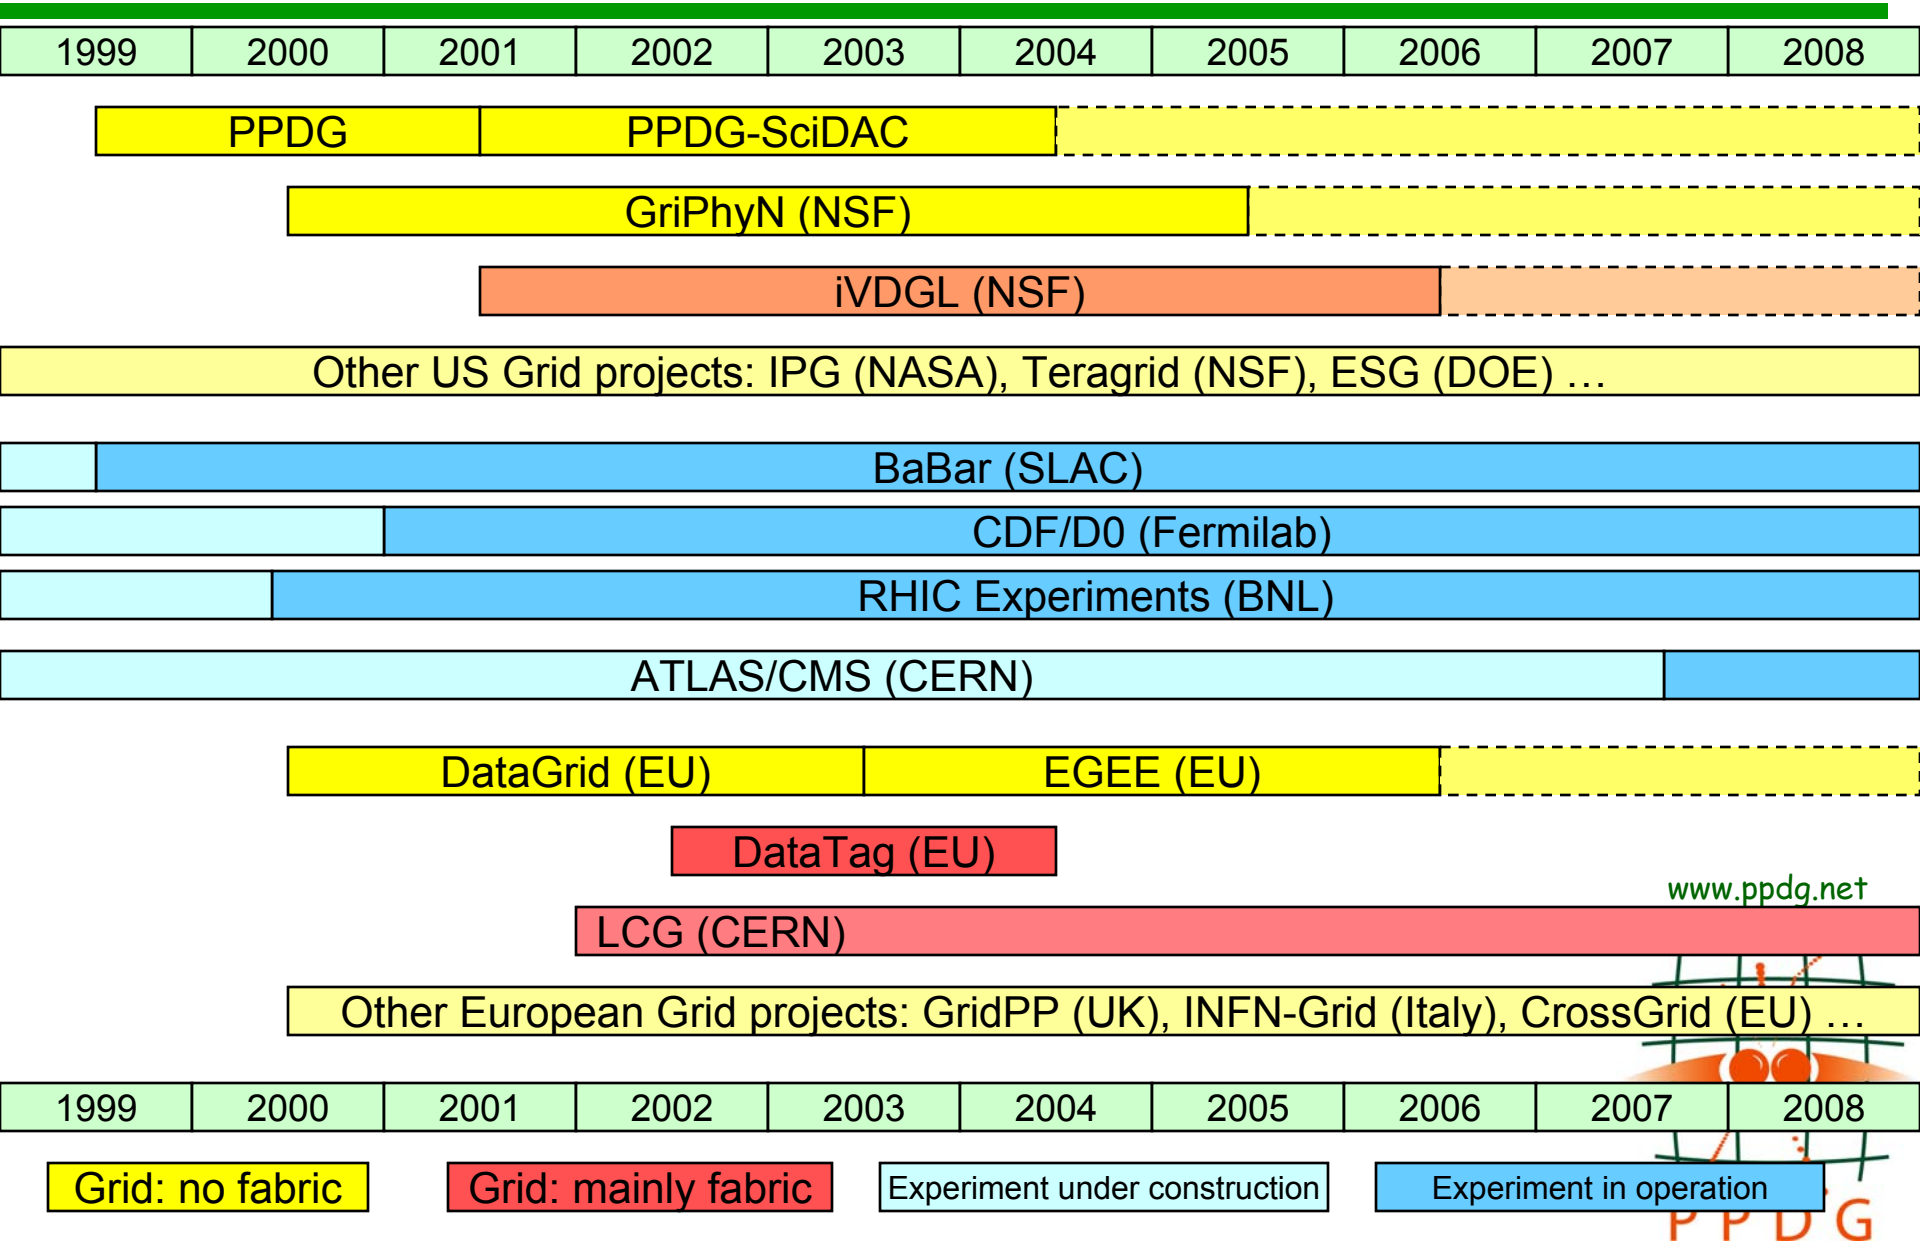

Coordination: HICB

HICB

HENP InterGrid Collaboration Board
(email list)

Meetings:

4 March 2001 : 23 June 2001

7 October 2001 : 17 February 2002

21 July 2002 : 13 Oct 2002

3rd March 2003

HIJTB

HENP InterGrid Joint Technical Board
(email list)

Meetings:

6 May 2002 : 8 April 2002

4 March 2002 : 18 February 2002

High Energy and Nuclear Physics InterGrid Web Pages

The HEPN Grid R&D projects (initially DataGrid, GriPhyN, and PPDG, as well as the national European Grid projects in UK, Italy, Netherlands and France) have agreed to coordinate their efforts to design, develop and deploy a consistent open source standards-based global Grid infrastructure. The guidelines for coordination and joint development by the projects are enunciated below. This collaborative effort can be referred to as HENP-INTERGRID .

HICB meetings are held three times a year co-located with the Global Grid Forum

We list s set of useful links reflecting some of the work of the board:

Collaboration Plan

Grid Laboratory Uniform Environment (glue)

glue schema

PPDG

GGF Research Group: Particle and Nuclear Physics Applications

Administrative: Group Name: **Particle and Nuclear Physics Applications Research Group**
Names and contact information for group chair(s): High Energy Physics Intergrid Coordination Joint Technical Board Chairs, Ian Brid (CERN, ian.bird@cern.ch), Ruth Pordes (fermilab, ruth@fnal.gov)
name and contact information for group secretary: Doug Olson (LBNL, dlolson@lbl.gov)
mailing list address: hijtb@hicb.org - will be augmented by pnpa@hicb.org
website (or web page) address: www.hicb.org/pnpa (currently www.hicb.org/jtb)

the following institutions and organizations will be working with and contributing to this research group:

Projects: EDG, LCG, CrossGrid, DataTAG, PPDG, GriPhyN, , PPARC, iVDGL, +++
Experiments: ALICE, ATLAS, BaBar, CMS, D0, STAR, TJNAF, BTEV, CDF, ++++
Institutions: ANL, BNL, Caltech, CERN, University of Chicago, University of Florida, Fermilab, INFN, KEK, LBNL, NIKHEF, SLAC, University of Texas at Arlington, University of Tokyo, Vanderbuilt, ++++

Charter for GGF Research Group on Particle and Nuclear Physics Applications (PNPA)

The Particle and Nuclear Physics Applications Research Group provides a forum for discussion of and issues related to particle and nuclear physics applications and production grids, with particular focus on the requirements from and activities of the community related to other GGF working groups and research groups.

The particle and nuclear physics community has over the past few years provided the driving applications and made significant contributorions to several projects (EDG, DataTAG, PPDG, GriPhyN, SiteAA iVDGL, Nordugrid) that have been actively involved in developing and deploying prototype and now production grid applications. Many of the participants in these projects are active participants or leaders of existing or new GGF working groups or research groups. Current areas of interest include security, high throughput data transport, data intensive management and access services, and information schema definition. As the community continues to develop and adopt higher level and more comprehensive services in can be anticipated that we will contribute to other areas of work such as production and operations services, scheduling and job description, databases, frameworks and applications, and network and resource management.

The PNPA Research Group will provide a venue to bring together domain application scientists together with the information technologists and researchers in grid technologies, at the leading forum for standards.

Specific topics of interest include:

www.ppdg.net

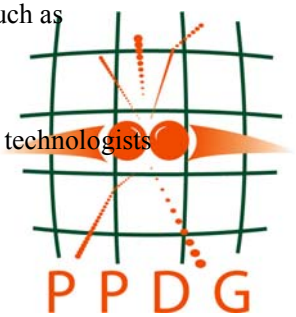

PPDG Summary

- **The first HE(N)P Grid Project**
- **CS – HENP collaboration**
- **Focus on end-to-end Grid deployment driving:**
  - Computer Science
  - Operational procedures (e.g. security)
  - Grid-middleware hardening, robustness and troubleshooting
- **A leader in the international coordination of Grid projects.**

www.ppdg.net

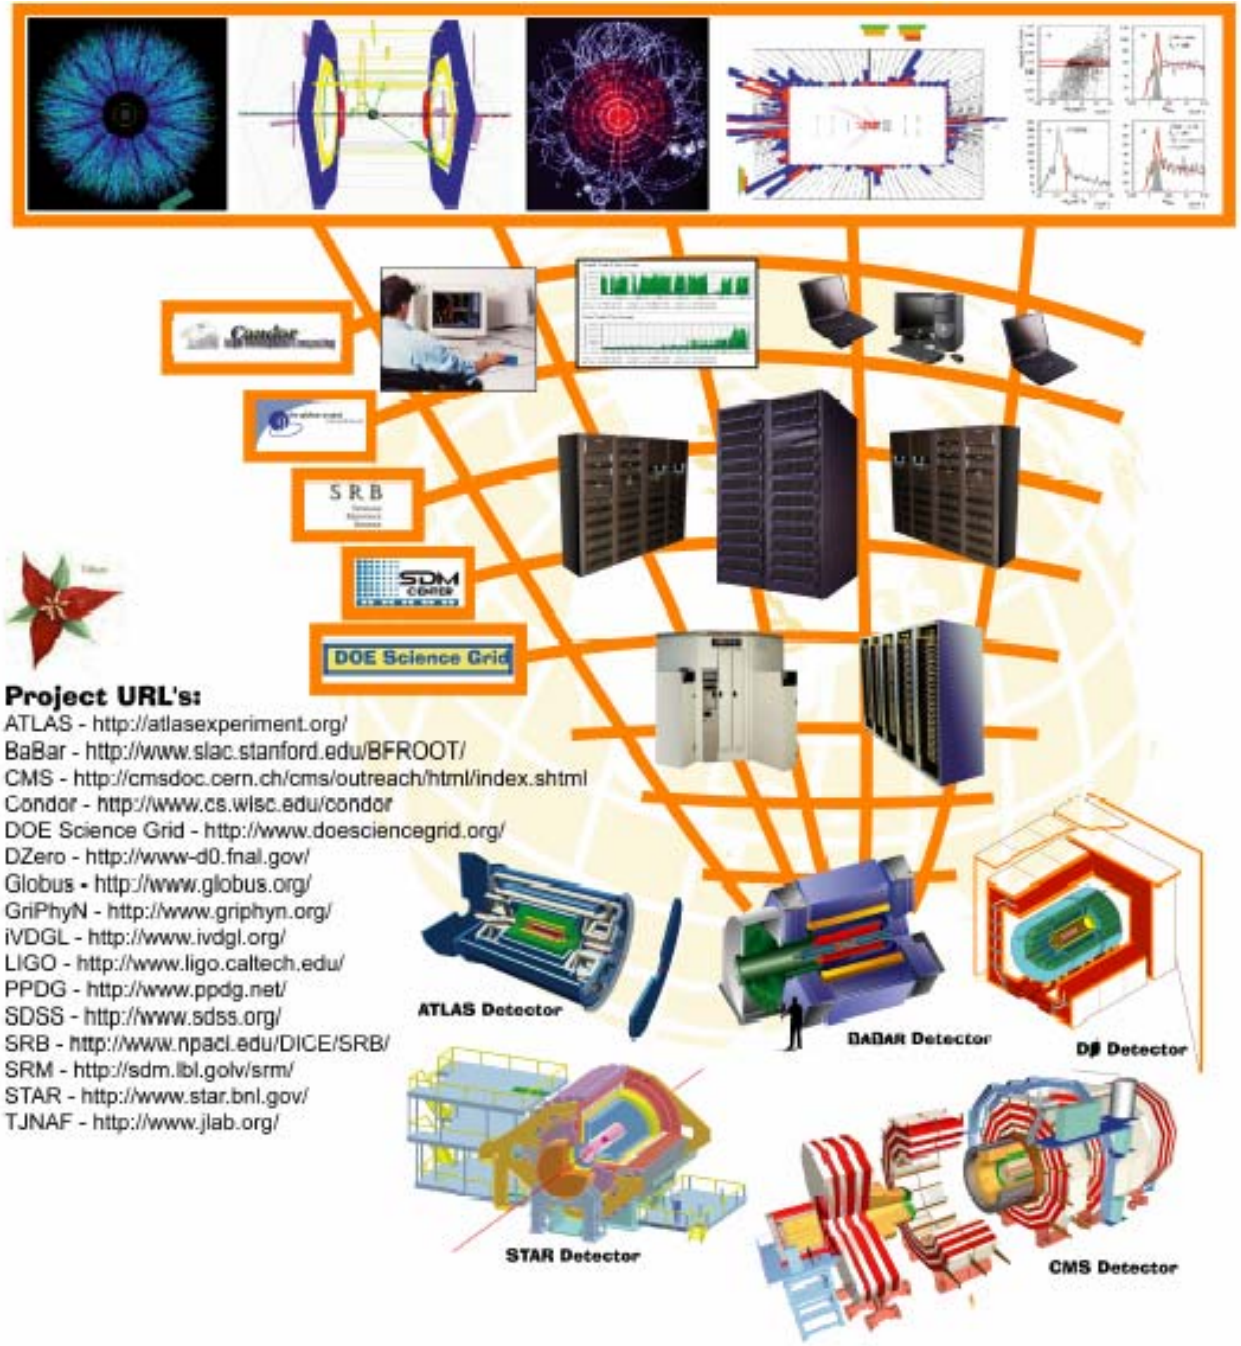

Project URL's:

ATLAS - http://atlasexperiment.org/
BaBar - http://www.slac.stanford.edu/BFROOT/
CMS - http://cmsdoc.cern.ch/cms/outreach/html/index.shtml
Condor - http://www.cs.wisc.edu/condor
DOE Science Grid - http://www.doesciencegrid.org/
DZero - http://www-d0.fnal.gov/
Globus - http://www.globus.org/
GriPhyN - http://www.griphyn.org/
iVDGL - http://www.ivdgl.org/
LIGO - http://www.ligo.caltech.edu/
PPDG - http://www.ppdg.net/
SDSS - http://www.sdss.org/
SRB - http://www.npaci.edu/DICE/SRB/
SRM - http://sdm.lbl.golv/srm/
STAR - http://www.star.bnl.gov/
TJNAF - http://www.jlab.org/

www.ppdg.net

PPDG