Computing Division

DOE Annual Review

March 19, 2003

Vicky White

Head, Computing Division

Outline

I. Mission, Activity Areas, and Organization

II. Activity Areas explained and accomplishments of the past year – highlights

III. Plans for next year and beyond

IV. Conclusions

Section I

Mission, Activity Areas, and Organization

Mission Statement – Nov 1, 2003

The Computing Division's mission is to play a full part in the mission of the laboratory and in particular

Participate in the Science

To proudly develop, innovate, and support excellent and forefront computing solutions and services, recognizing the essential role of cooperation and respect in all interactions between ourselves and with the people and organizations that we work with and serve.

Serve the Program

Collaborate

Drive the Program

Five Activity Areas

1. Provide services, tools, and components, and operate computing facilities that serve the lab and the scientific program broadly.
2. Provide dedicated help, leadership and active participation in running and approved experiments and other lab scientific programs (including support and expert help to the Beams Division).
3. Work on projects funded competitively outside the base budget – e.g. SciDAC & GRID projects.
4. Participate in planning and R&D for future experiments/lab activities.
5. Run a computing organization and computer center.

How are we organized for these 5 activity areas?

- Division Re-organized Dec 1, 2002
  - Started the job of Division Head Nov 1, 2002
  - Appointed 3 Associate Heads – one each for Operations, Planning and Projects/Proposals
  - Bob Tschirhart joined division as my deputy March 1, 2003 – focus on the Scientific Program
  - Went from 16 "boxes" to 9 to consolidate functions and empower staffing changes/evolution

New Computing Division Organization

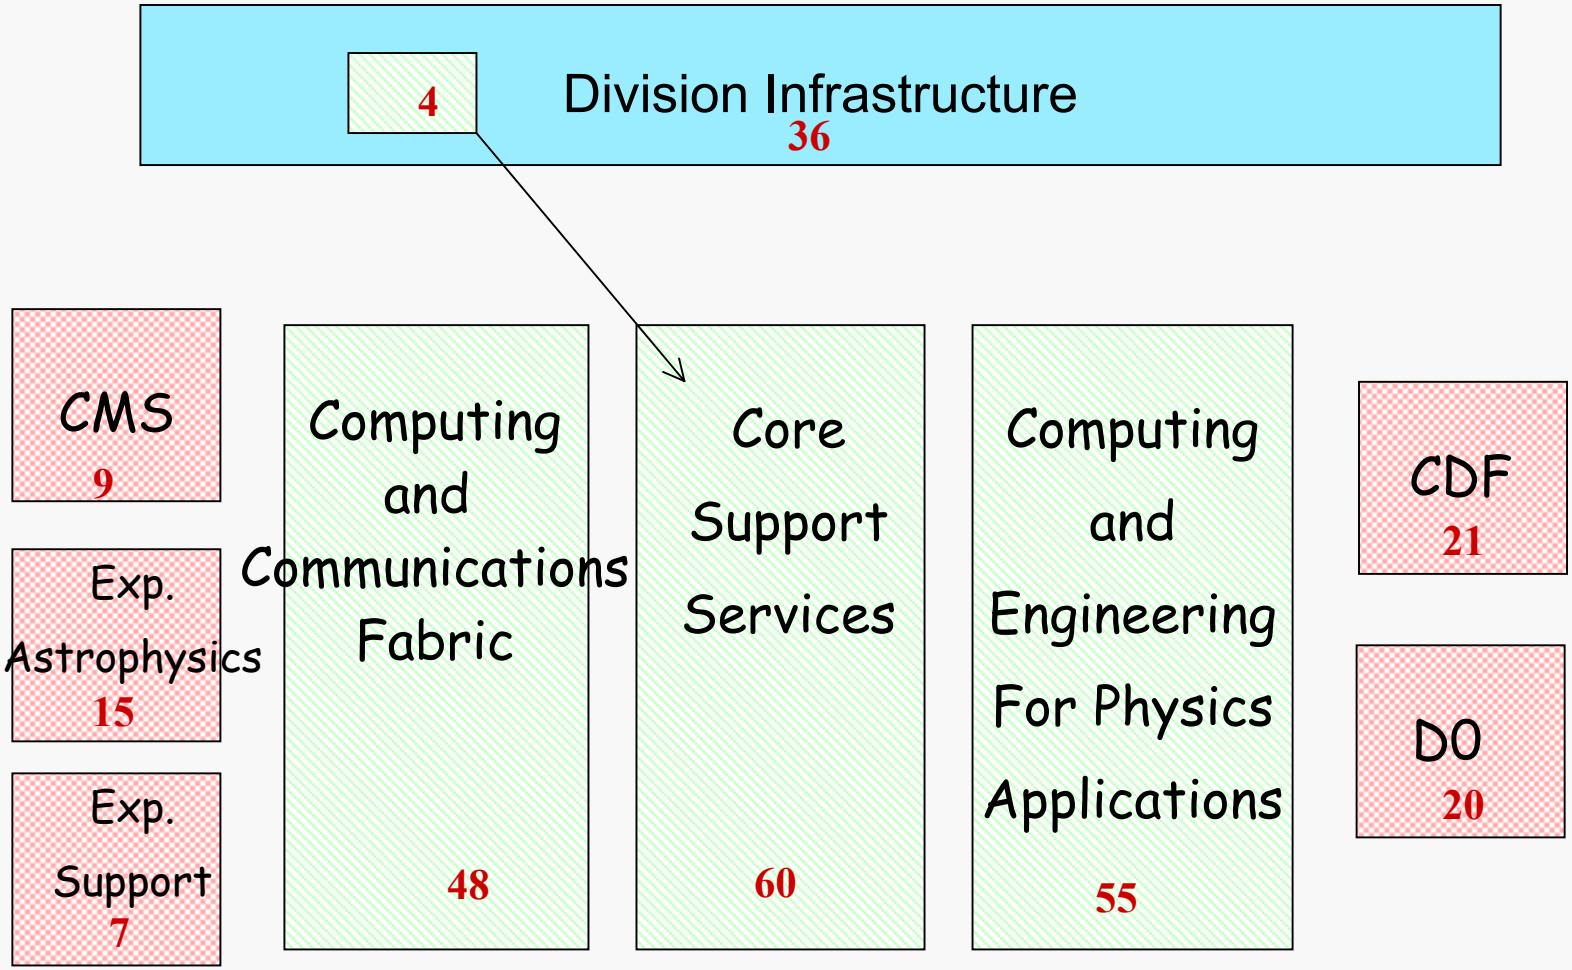

Scientists of all sorts, Engineers, Technical, Computing, Admin = 275

High Level Summary of Titles - September 2002

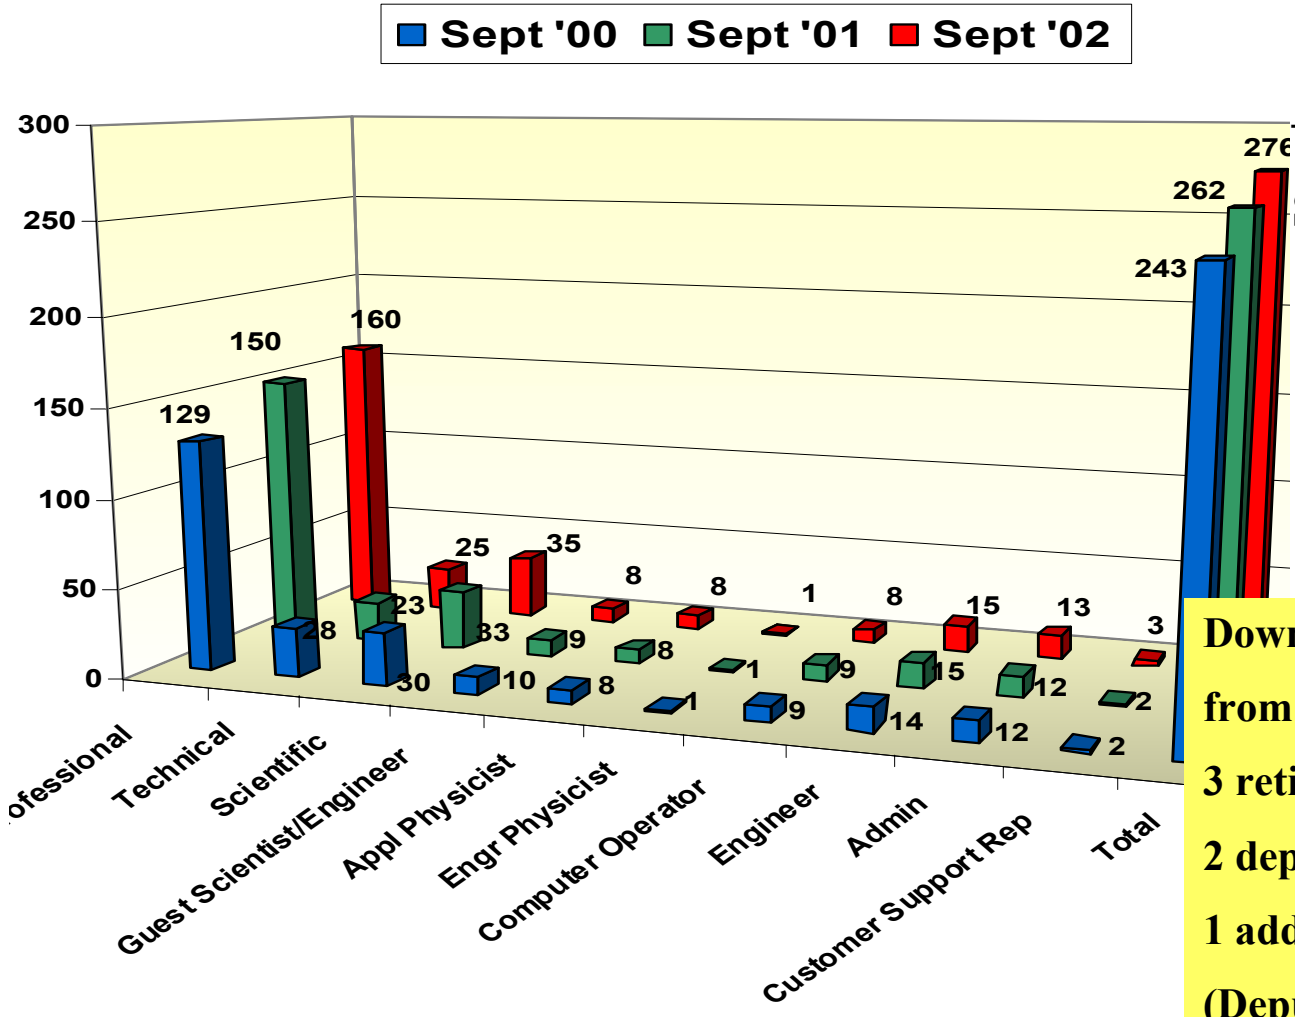

+ 3 Guest Scientists = 279

Down 4 people from Sep2002

3 retirements

2 departures

1 addition (Deputy Head)

Section II

Activity areas explained and Highlights of the Past Year

(1) Services, Tools, Components and Computing Facilities – used by all

- Campus Network and Wide Area Network connectivity (17 FTEs)
- Lead the lab's Cyber Security program (4FTEs)
- Operate central mail, print, file servers (3 FTEs)
- Lead Windows activities and services (11 FTEs)
- Helpdesk (4FTEs)
- Manage hardware and software maintenance contracts & vendor services (3FTEs)
- Database administration and database applications (lot for Run II) (15FTEs)

(1) Services, Tools, Components and Computing Facilities – used broadly

- Provide Engineering and technical assistance to running experiments (10 (+3) FTEs)
- Run PREP (Physics Research Equipment Pool) – modules and computers, repairs,technical help (14 (-3)FTEs)
- Build & operate central Data Storage and data handling systems used by all experiments, theory, SDSS (10FTEs)

(1) Services, Tools, Components and Computing Facilities – used broadly

- Procure and manage compute resources (all Farms)
- + Operate and evolve shared analysis machines and lab's AFS (distributed file system) (15 FTEs)
- Operating systems and tools (Linux) (2FTEs)
- Develop & support Physics Analysis Software used HEP-wide (e.g. ROOT, GEANT, Generators, C++ libs, graphics, Compilers,Build tools, etc..) (9FTEs)
- DAQ and online tools of broad utility (2FTEs)

(1) Services, Tools, Components, Facilities used broadly

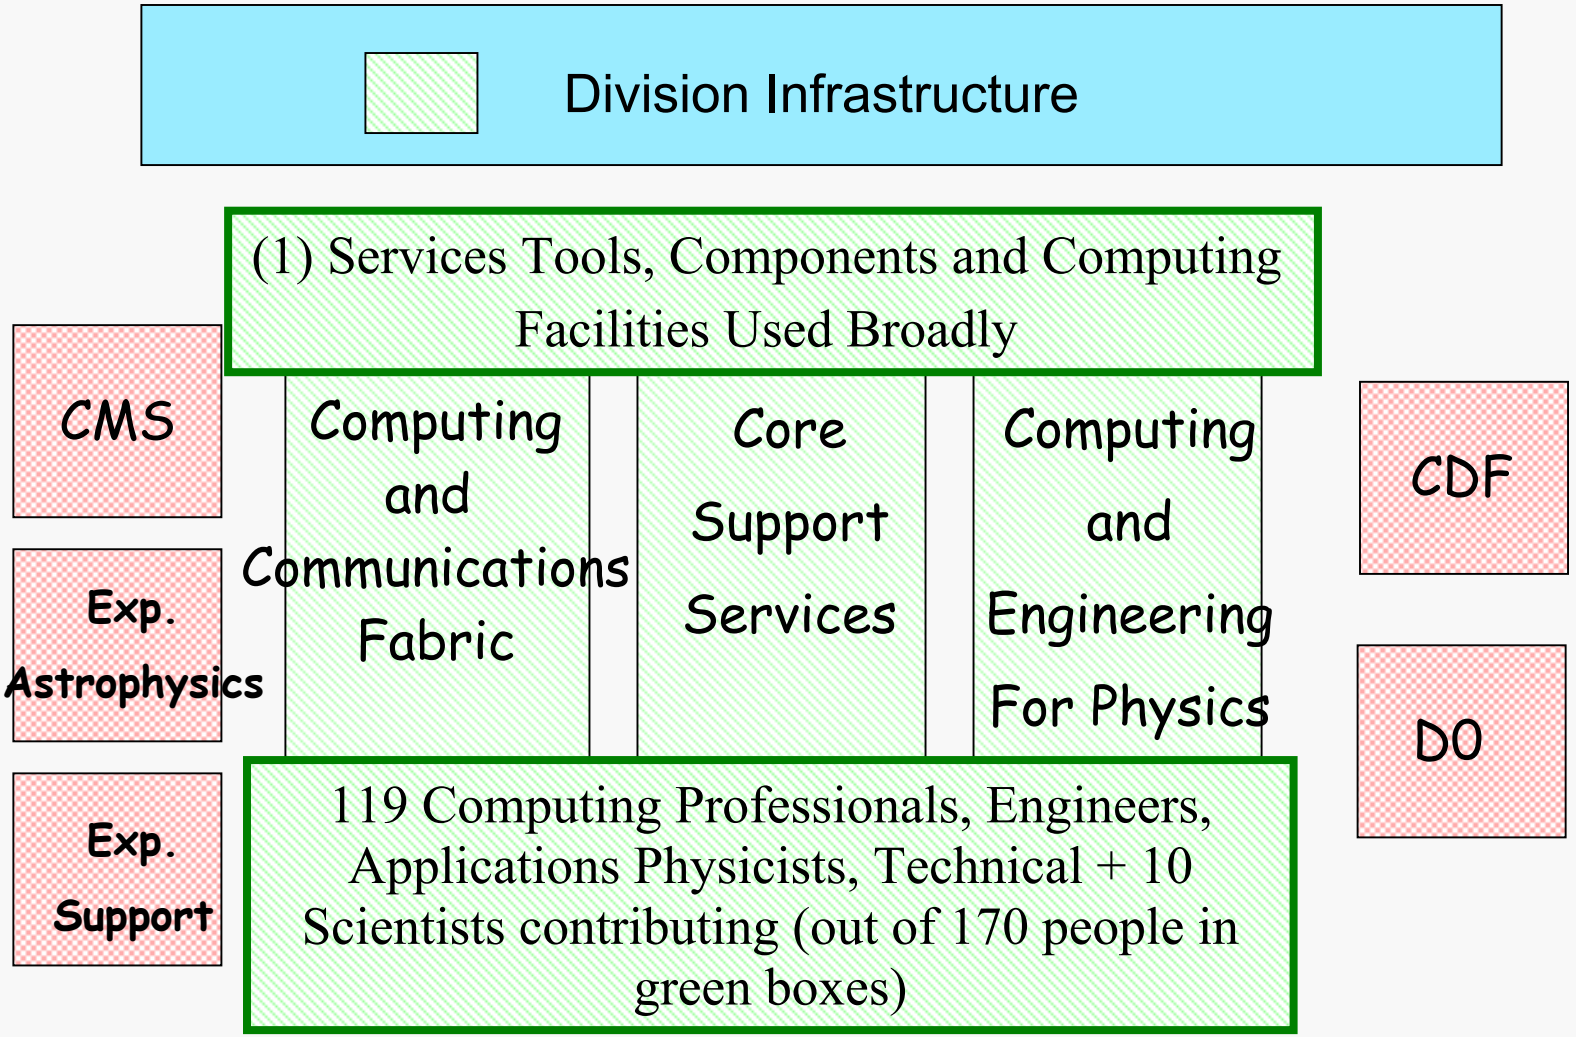

Highlights

Engineering Projects for CDF and D0

Successfully completed engineering and system integration projects for CDF and D0 Run IIa

- D0 Level 3 Trigger – Lead role in replacement of Level 3 Trigger (hardware/software integration project)
- 'SVX Modules' used in SVX DAQ, SVT Trigger and D0 DAQ
  - Six full module designs since 1994 with 15 versions and many test modules.
  - ~1000 modules delivered to the experiments. Ongoing support.
- D0 projects
  - Clock and Timing distribution system
  - CFT mixer system to reorder mixer streams for the trigger
  - Alpha L2 trigger processor repair and support
  - Alpha L2 trigger processor bus interface
  - L1 Calorimetry trigger daughter card
  - Analog Front End Board layout

Centrally managed Storage Systems used by entire program

- 4 Robot installations (STK and ADIC)
- 77 Tape drives; 5 different technologies (M2, 9840, 9940A, 9940B, LTO)
- 13,600 tapes
- 673 TB of data on tape

- dcache (collaboration with DESY) disk cache in use by CDF, CMS and MINOS

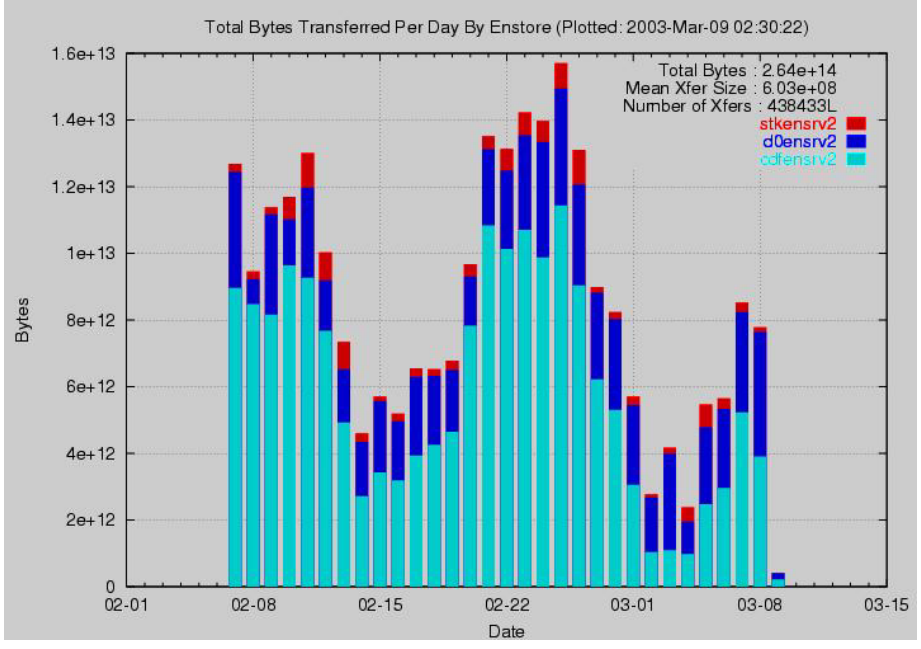

Peak of 15 TB of data per day moved in/out of Storage systems !!

Networks

- EsNet network connectivity increased to OC12 - i.e. 622Mbps (from 155Mbps) in October 2002
- Networks are essential for data intensive commodity computing

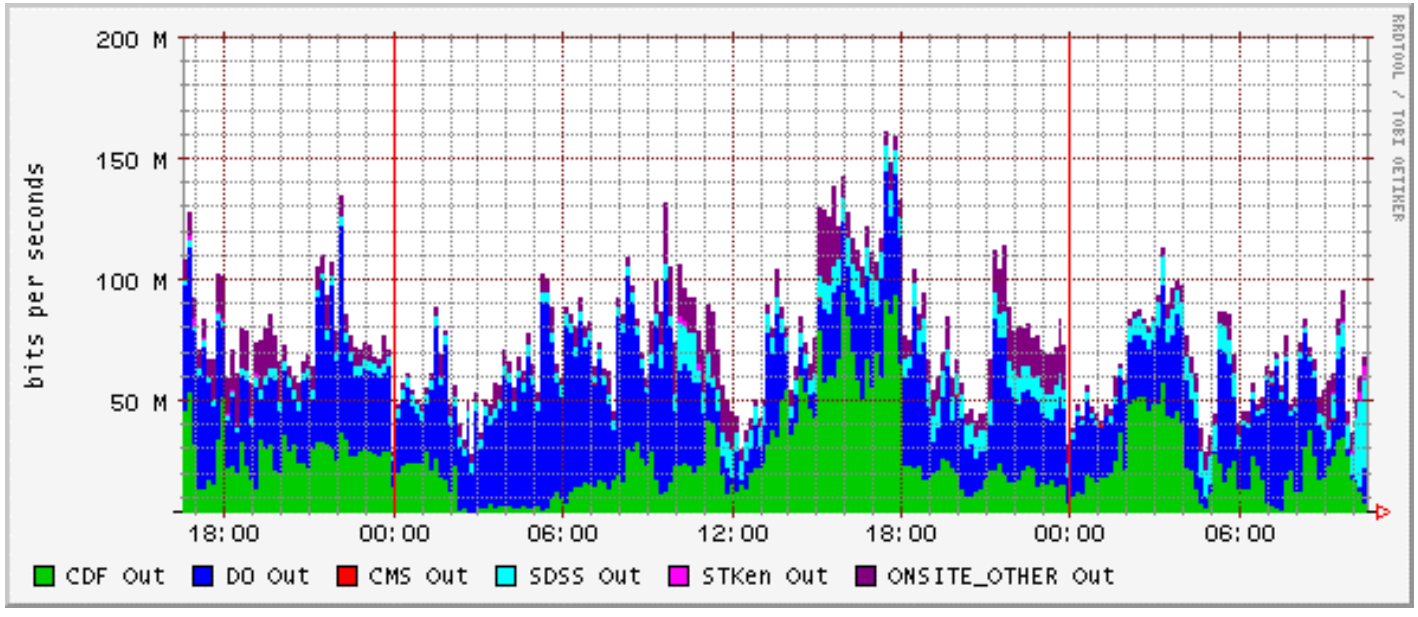

Farms (Production Processing resources)

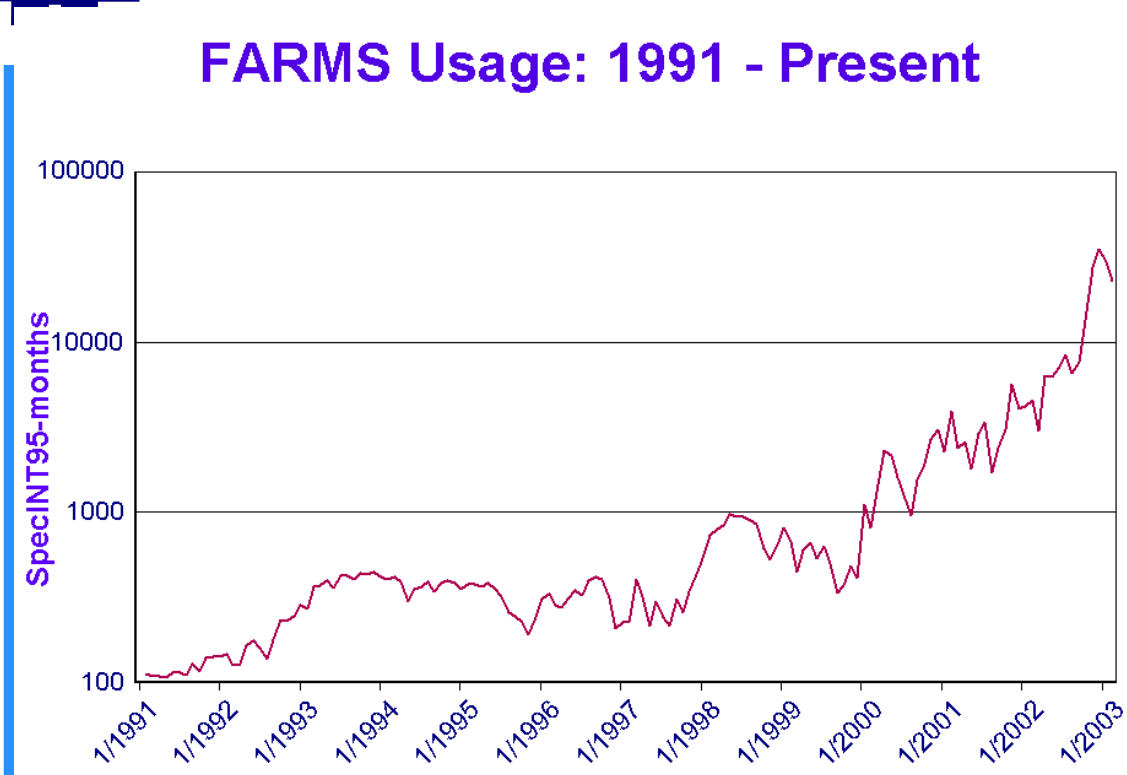

CMS 1U

←Fixed Target →

← Run I →

← Run II →

CDF and D0 Farm Production

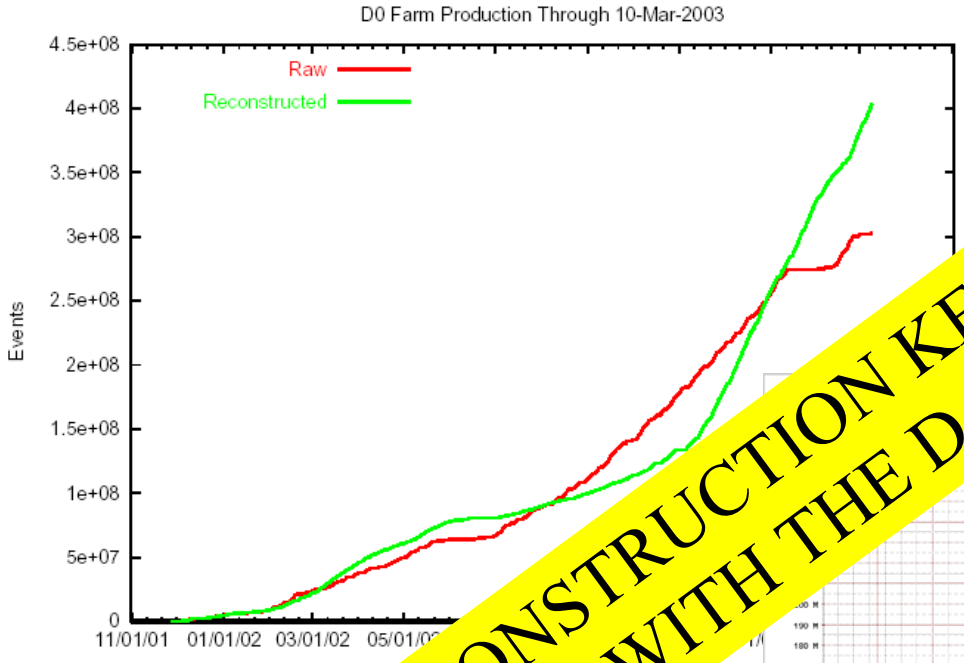

D0 Reconstru

RECONSTRUCTION KEEPING UP WITH THE DATA

CDF Reconstruction

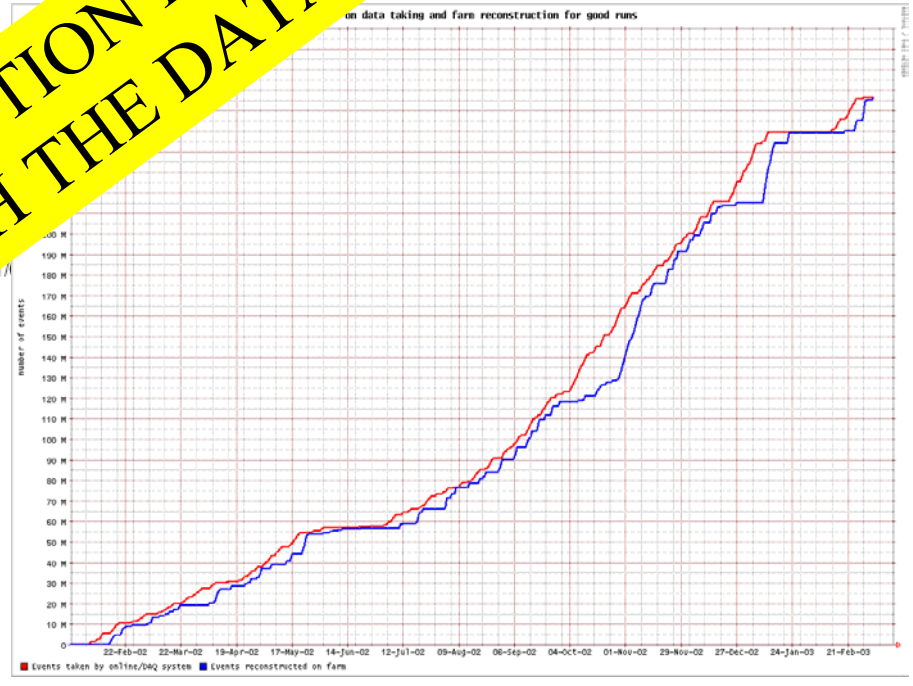

Lab-wide Windows policy group

- Windows Policy Working Group formed
- Migration of Windows systems to Windows 2000 Domain almost complete
- Good cooperation and coordination site-wide

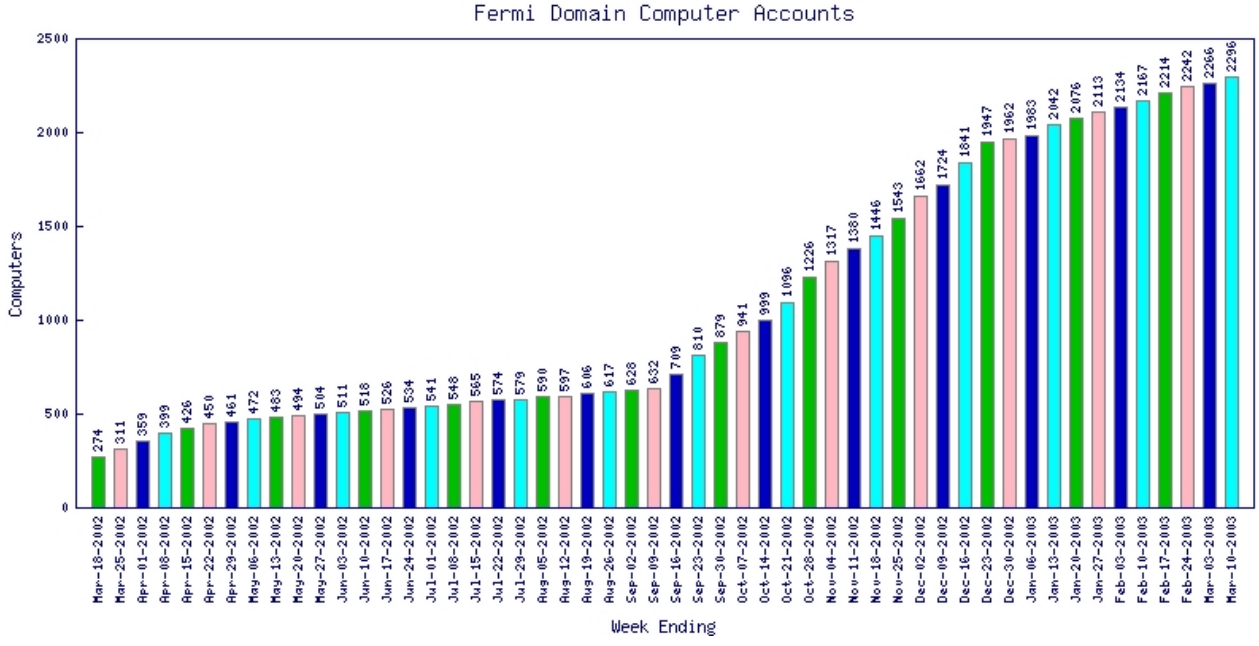

(2) Dedicated help and active participation in science program

	FTEs(non-scientist)	Research Scientists	
Accelerator	9	3	+ SciDAC Accel
CDF	10	11	
D0	8	12	+Grid Projects
Run II	5		
CMS	8	3	+Grid Projects
CMS Detector	1.5		
NUMI/MINOS	2	2	
SDSS	4	8	+Grid Projects
Theory		2	+SciDAC Lattice
CDMS	0.3		
miniBOONE		1	
Education	1.5		
	51.3	**43**	

(2) Dedicated help and active participation in science program

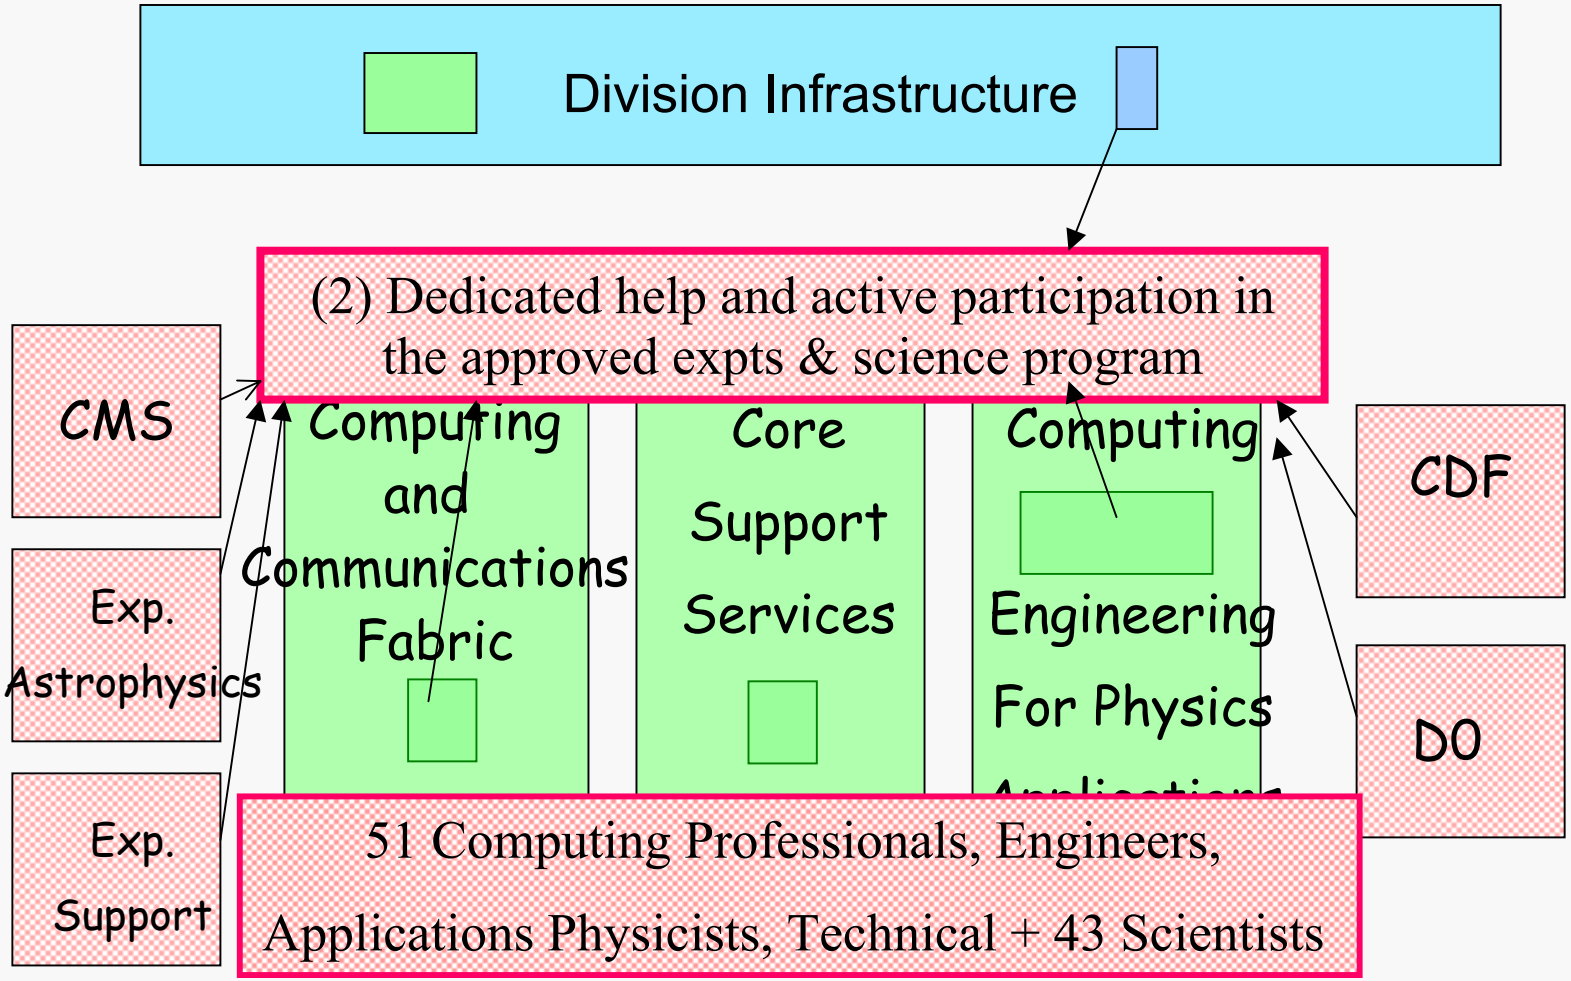

Highlights

Beams Division help

- 11 FTEs of help on projects from throughout the division. 3 scientists strongly involved.
  - Patty McBride leading the effort
  - Paul Lebrun immersed and lauded for his effective help
  - Pushpa Bhat involved in Review organization
- We DO have expertise to offer
  - can bring tools and experiences from experiment DAQ, instrumentation, software development and data analysis, migration projects, project management?
- HIGHLIGHTS
  - Shot Data Analysis work
  - Java tools work
  - Instrumentation & technical for BPMs and Flying Wires

20 people involved – limited by nature of tasks available

CDF Computing

- Central Analysis Farms (CAF)
  - March 2002: CAF did not exist. Prototype with 16 CPUs & 2 TB disk.
  - July 2002: CAF for Summer 02 conferences. 120 CPUs, 30 TB.
  - March 2003: CAF meets needs for Winter 03 conferences. 600 CPUs (1THz), 180 TB disk.
    - Many hardware and software bumps along the road – with IDE disks, controllers, dcache. State-of-the-art system!
- Data Handling
  - Archive: 250 TB of data in robot March '03 (100 TB in March '02)
    - Reused 120 TB of old production data tapes
    - Users read data from disk at 200 MB/s in Mar '03, up from ~20 MB/s in '02.

SAM Data Handling System at DO

Integrated Files Consumed vs Month (DØ)

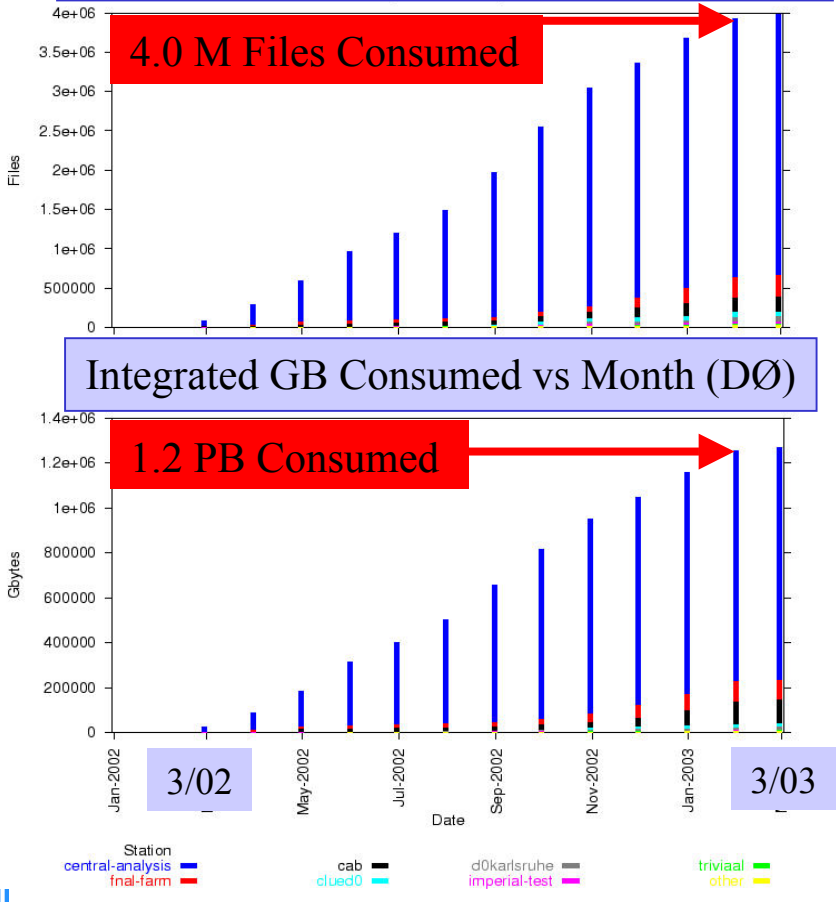

Summary of Resources (DØ)

Registered Users	600
Number of Stations	56
Registered Nodes	900
Total Disk Cache	40 TB
Number of Files	1.5M

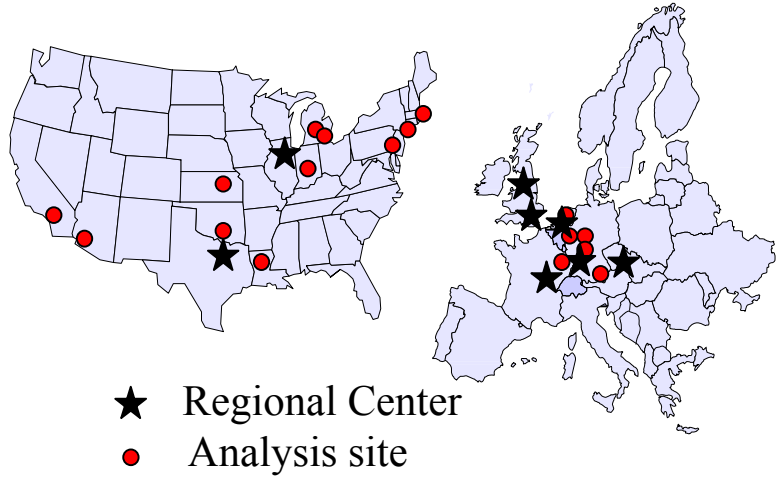

CMS Software and Computing Project

- Fermilab is the host lab of U.S. CMS
- Fermilab Computing Division (CD) hosts the U.S. CMS Software and Computing Project.
  - Management of project is in CD - L1 project manager, Lothar Bauerdick
  - L2 projects
    - User Facilities (Tier 1 and Tier 2 centers)
    - Core Application Software
  - Ian Fisk will be joining CD in April as User Facilities L2 project manager
  - Project ramping up – more slowly than planned, due to lack of funding. Difficult to manage a project with a changing funding profile. U.S. Leadership requires people in place early.
  - Working closely with University partners

US-CMS Integration Grid Testbed

- Fermilab (Tier1)
  - 40 dual 0.750 GHz processor machines
- Caltech (Tier2)
  - 20 dual 0.800 GHz processor machines
  - 20 dual 2.4 GHz processor machines
- San Diego (Tier2)
  - 20 dual 0.800 GHz processor machines
  - 20 dual 2.4 GHz processor machines
- Florida (Tier2)
  - 40 dual 1 GHz processor machines
- CERN (LCG Tier0 site)
  - 36 dual 2.4 GHz processor machines

- Total:
  - 240 0.85 GHz processors: Red Hat 6
  - 152 2.4 GHz processors: Red Hat 7

1M events fully simulated and reconstructed on the IGT

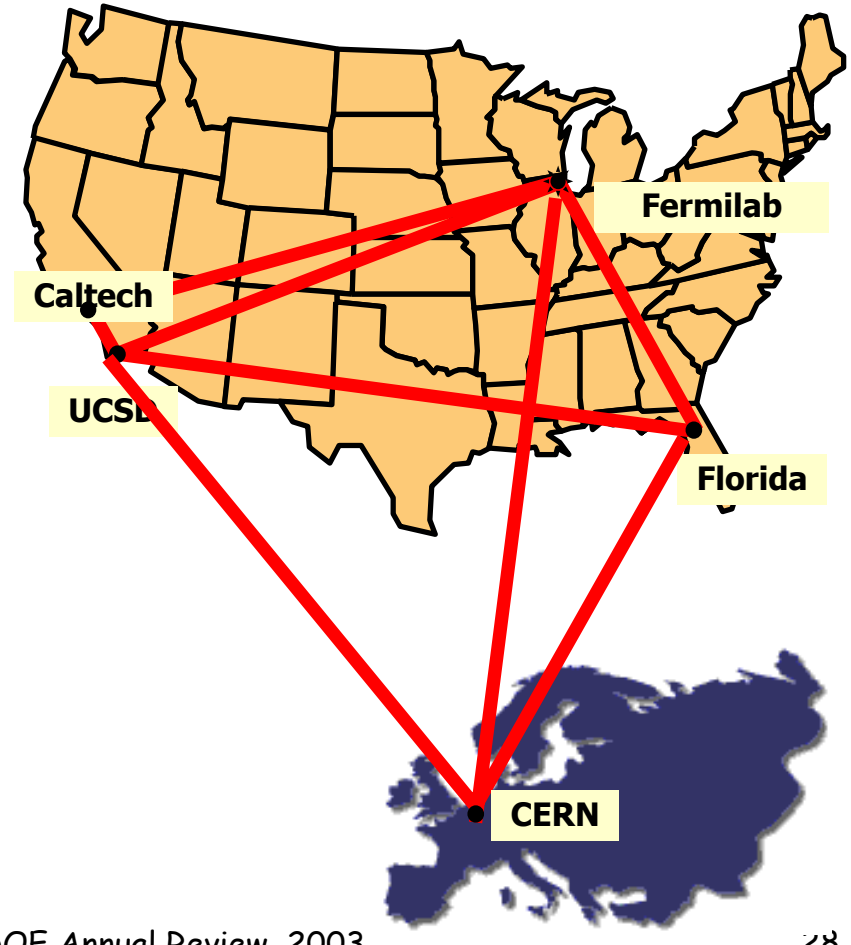

Participation in Physics Analysis

- CDF
  - 2 physics group convenors – Stephan Lammel and Pasha Murat,
  - 2 "wine&cheese" talks at Fermilab
- D0 – increased participation in physics groups
  - 1 physics group convenor – Gustaaf Brooijmans
  - 3 winter conference talks
- FOCUS, KTeV, NuTeV
  - 1 KTeV winter conference talk
- SDSS – several new results including discovery of a distant ring of stars found circling the Milky Way – Brian Yanny (CD) one of primary authors with Heidi Newberg (formerly CD).

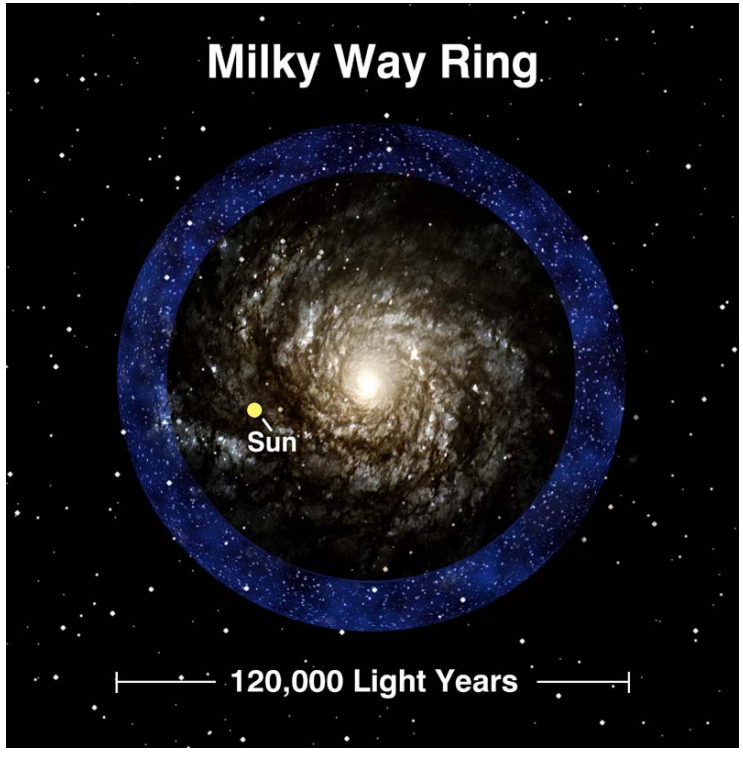

(3) Participation in Competitively funded Proposals & special Projects

- KDI (finished) (NSF) — funded 2000-2002
- SciDAC: PPDG (DOE)
  - PPDG_1 — funded 1999-2000
  - PPDG_2 — funded 2001-2003
- GriPhyN (NSF) — funded for 2000-2004
- iVDGL (NSF) — funded for 2002-2004
- NVO (NSF) — funded for 2002-2006
- SciDAC: SRM (DOE) — funded 2001-2003
- SciDAC: Lattice QCD (DOE) — funded 2001-2003
- SciDAC: Accel. Simulation(DOE) — funded 2001-2003
- RTES – BTeV (NSF) — funded 2002-2007
- HEPIC and DOE web page support — ongoing
- Network Monitoring (w/SLAC) (DOE) — ongoing
- DOE Computing Advisor (DOE) — 2003
- ARC funding for SDSS — ongoing

Total of $1.6M in FY03

Highlights

GRID Projects

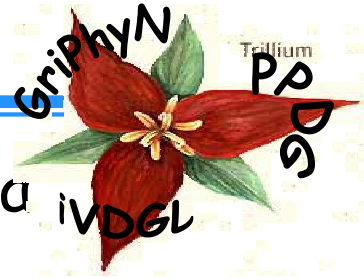

- Fermilab (Ruth Pordes in particular) playing a lead role in bringing together and coordinating US HEP Grid Project efforts as "Trillium" confederation of PPDG, iVDGL, GriPhyN GRID projects.
- Also much work under auspices of Joint Technical Board of HENP International Coordinating Board (HICB).
- Progress in coordination of LHC Computing Grid, US-CMS and European Data Grid, and Run II SAM-GRID.
- SC2002 demonstrations (WorldGrid).

Lattice QCD Facility Operational

- 80-node Pentium-III cluster in production since January 2001
- 48-node SciDAC Xeon cluster in production since July 2002
- 128-node SciDAC Xeon cluster just brought into production (late January)
  - **Top500.org score (High Performance Linpack) is about 570 GFlop/sec**
  - **this ranks the cluster near #96 on the list**
  - sustained performance for MILC Improved Staggered is about 60 GFlop/sec
  - at least 10X better than ACPMAPS
- Software contributions:
  - FNAL SSE/SSE2 matrix algebra codes now part of MILC, will port to QLA

Lattice QCD

Muon Lab

Complex Interconnects

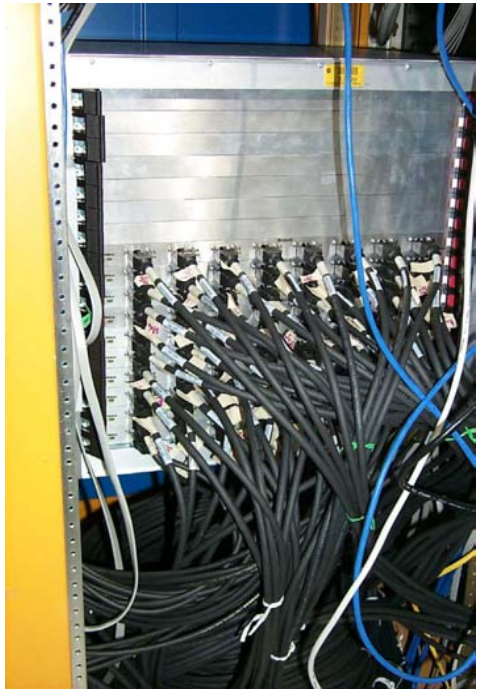

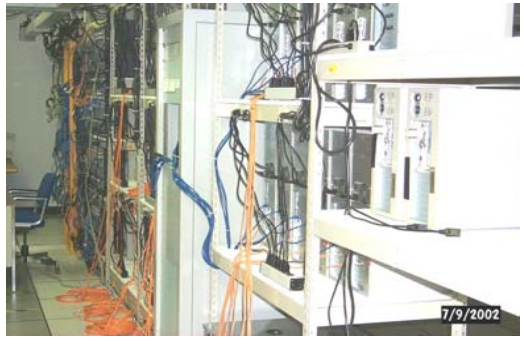

Room To Expand

128 Nodes

Lattice QCD "Commodity" Supercomputer

Highlights – Accelerator Simulation

- Multi-institution collaboration to develop the next generation of beam dynamics modeling tools (and other Accel. Simulation tools)
  - At Fermilab – Booster Simulation (Panagiotis Spentzouris and Jim Amundson)

IPM is the only Booster device with turn by turn resolution ⇒ needs to be calibrated

Use "flying beam" technique and MWPC to obtain IPM calibration!

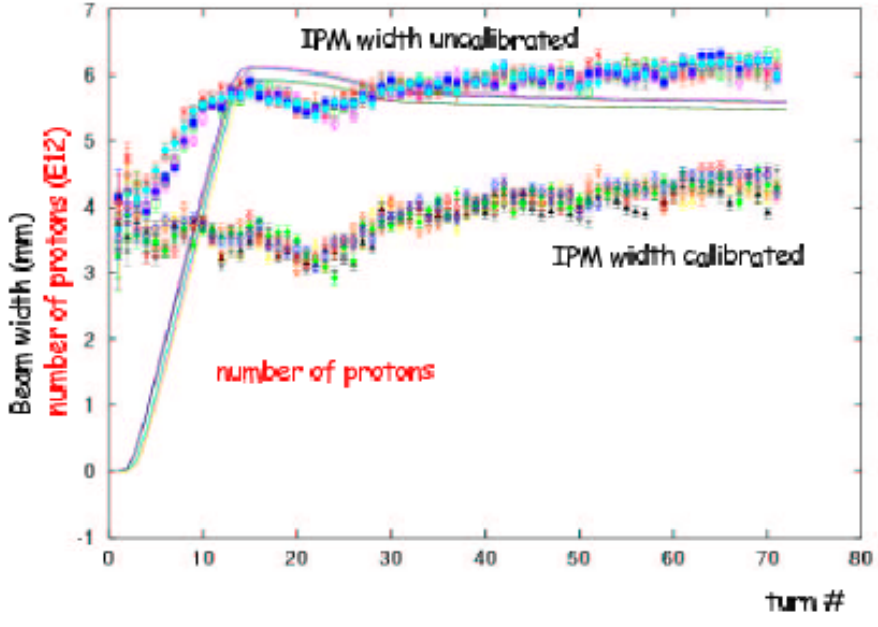

Booster model validation

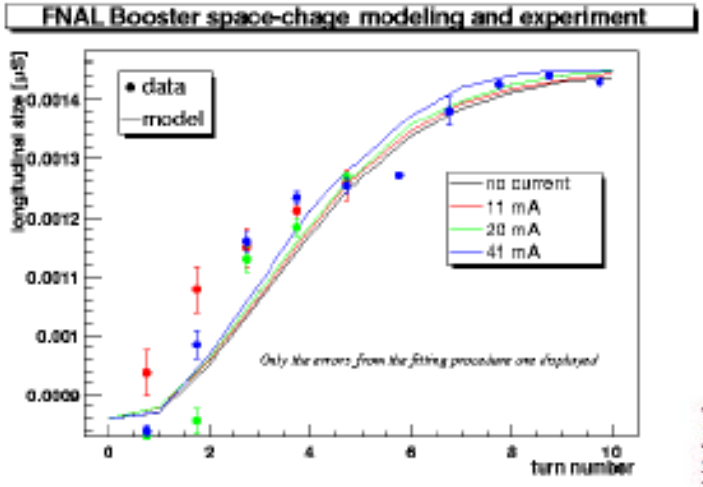

Model validation effort:

Data/MC comparison, a rare phenomenon in Beams Physics!

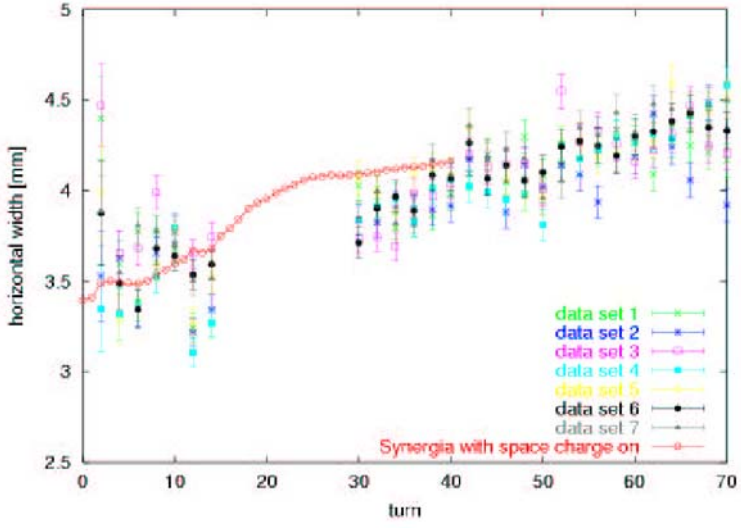

More studies underway in collaboration with BD

(4) Planning and R&D for future experiments & lab science activities

	FTEs (non Scientists)	Active Scientists
BTeV (DAQ &Pixel R&D)	8	
BTeV Comp & Trigger	1.8	1 (+3 -> Beams help)
CKM (DAQ)	1	2
Neutrino Factory and Linear Collider	0	0
	10.8	3 (+3)

(5) Running a Computing Organization and a Computer Center

Planning	Steve Wolbers
Budget	
Administrative	
Operations & Ops metrics	Gerry Bellendir
Computer Center Building & Space management	
Safety	
Computer Center Operations	
Proposals, Education & Outreach	Ruth Pordes
Project Organization and Tracking	

Summary of Computing Division Effort

	Activity	FTE's (Non scientists)	Scientists
1	Services, Facilities used broadly	119	(10)
2	Dedicated help & participation in approved science program	51	43
3	SciDAC & other projects	15	(6)
4	Planning, R&D for future expts	11	3 (3)
5	Running division & comp.center	34	(2)
		230	46

Section III

Plans for the future

Run II

- Planned investment of $18M now in place, operations started in 02
  - Need replacement of ~20-25% of disk&cpu /year
  - Need for $2M/year/expt validated by Run II review
  - This year we will spend $1.45M/expt
- Migration from SGI servers to Linux Analysis Facilities is on a fast track in FY03. Complete in FY04.
- GRIDs and Distributed Computing for Run II are a reality/necessity
  - Convergence of tools and approaches between Run II and LHC is in view. Fermilab will continue to play a lead role in this.

Wide Area networking

- Fiber to Starlight this year (ComEd)
- Strategic direction –
  - Overflow bulk traffic to Universities
  - Connection to CERN, Nikhef, and other Eu sites
- Very important for Run II and CMS
- Opportunities to cost share
- Collaborate on new approaches – testbed => production

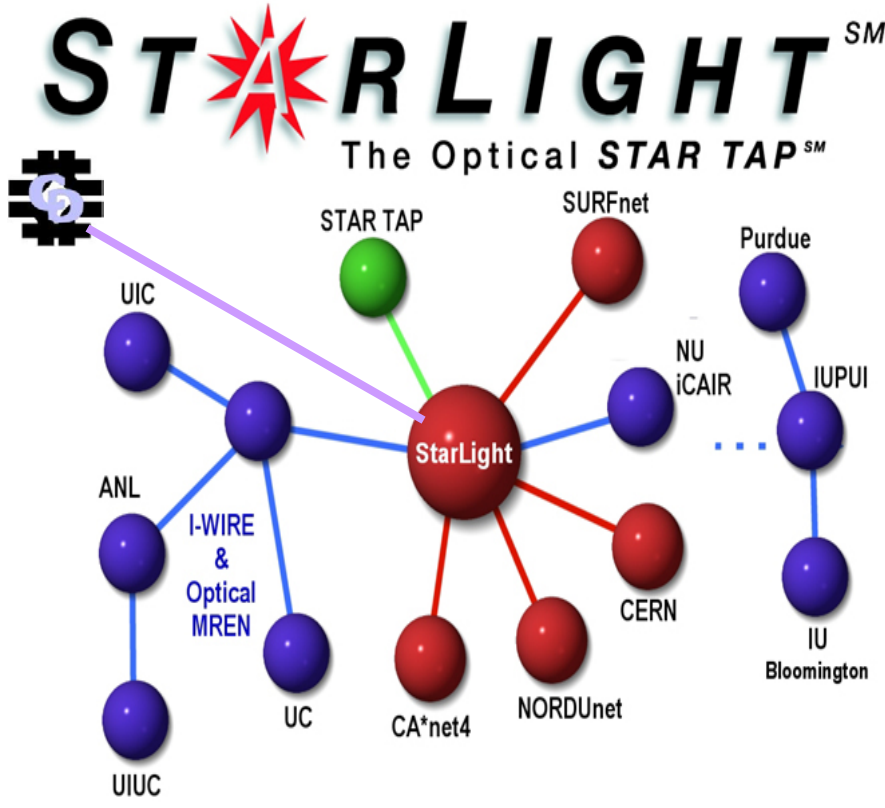

Starlight: optical networking interconnection point downtown Chicago (*710 Lake Shore drive*)

Owned by Northwestern University

US-CMS Software & Computing at Fermilab

- In FY2003 at Fermilab:
  - Equipment budget: $650k => make a significant Tier 1
  - Personnel ramps by 6.5 FTE
    - Will fill by internal transfers during FY03
- Much leverage of (and alignment with) CD services and facilities
- Strong leadership and management of the project at Fermilab. Play big role in Data Challenges.
- Continue to work out roles/ways of working with the LHC Computing Grid Project
  - US CMS representative on SC-2 committee – Lothar B.
  - US Representative on Grid Deployment Board - Vicky
- Coordination of ITR proposals for US-CMS
  - Joint proposals with ATLAS, others

Other Plans for the Future

- Cyber Security – increase effort
- Network – zoned architecture, redundancy
- Enterprise-wide backups
- Making our Computer Center Fabric – Grid enabled
- Databases – solve Oracle licensing issue and/or move to free database
- Division-wide project on Analysis Farms and File servers (Michael Ernst to lead this)
- Settle on portfolio of Run IIb, BTeV, CKM Engineering projects – and get going
- Ever more worldwide and cross-disciplinary collaborative efforts with labs and universities

FCC power and cooling

- Run2 and CMS require massive PC computing clusters
  - Very high physical density
  - ~200 Watts per CPU chip (similar for SMP and PC)
- Will hit FCC infrastructure limits in FY03
  - Working on options
- Carrying out Engineering Study of Options for FY03 – FY05
- Long term: substantial infrastructure overhaul or new building

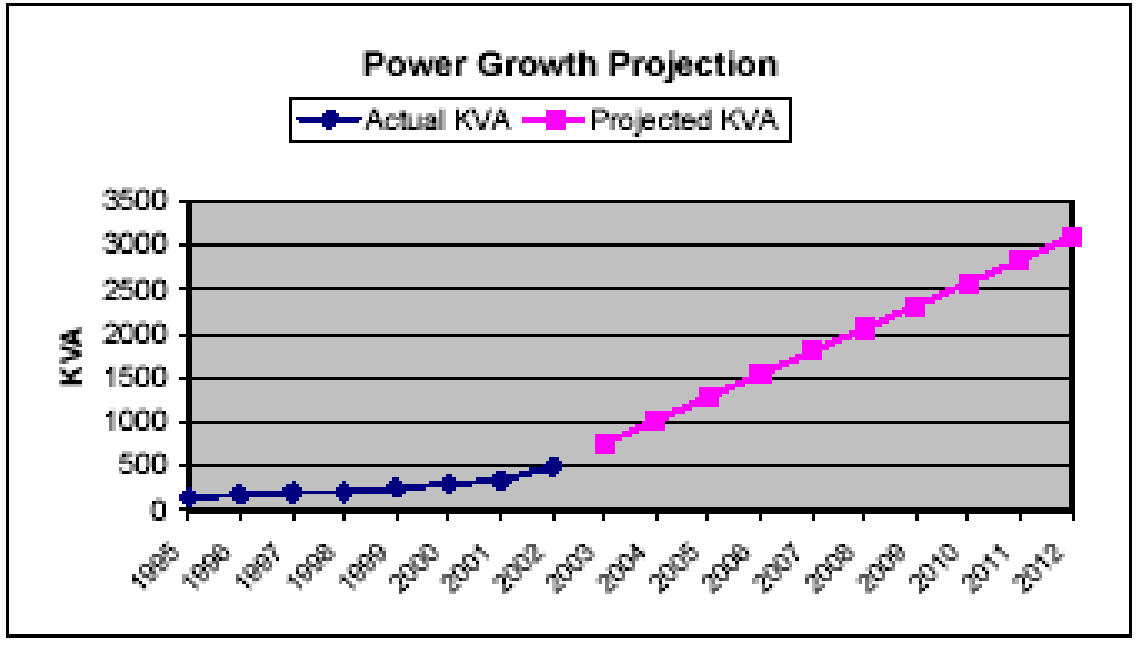

Section IV

Conclusions

Conclusions

- Highly leveraged and efficient common services, working in collaboration with partners to provide leading-edge solutions.
- Strong and active participation in the Physics Program of the lab and in direct support for it, including new and successful involvement with the Beams division.
- US-CMS Software and Computing, led by Fermilab team, playing leading role in CMS and in LHC Computing Grid Project.
- Fermilab Computing is in an excellent position to support the physics program of the laboratory and to participate in and foster partnerships with our collaborators, computer scientists, universities.

Conclusions

- By participating in these collaborative projects and by building working relationships we are able to extend our capabilities in:
  - Distributed and GRID computing (important for most of our collaborations)
  - Lattice QCD
  - Beam simulations
- Many challenges for the future – in technology, engineering, large collaborative project management, physical plant and staff evolution, and budget.
  - People are the key. Great staff in CD!