CDF & D0 Computing

Version 0.1

2 / 7 / 2008

GP Yeh

Thanks

Amber's Starting List

Amber's presentation
www-hep.uta.edu/~d0race/d0rac-wg/amber-future-budget.ppt

Andy gave me Lucchesi/Snider "Offline status and plans"
12/2007 presentation to CDF Executive Board
http://www-cdf.fnal.gov/internal/WebTalks/
http://hcp2006.phy.duke.edu/HCP2006-science.html
D0 Computing and Analysis Model by Amber
CDF Computing and Analysis Model by Pierre Savard

Jason Allen gave me 2 useful D0/CDF links:
http://rexganglia1.fnal.gov
http://d0om.fnal.gov/d0admin/faultlog/

Roman Lysak
enstore (tape usage at CDF/D0) + dCache
http://www-ccf.fnal.gov/enstore/
network:
fndcg0.fnal.gov/~netadmin/nwm/cgi-bin/temp/core.html
local CDF CAF and CDF farms outside Fermilab:
http://cdfcaf.fnal.gov/

Paris CDF Week Collaboration Meeting
http://lpnhe-cdf.in2p3.fr/cdf_parismeet/

Roser	
Glenzinski	Preparing for summer conferences
Hahn	Detector Operations Status
Lucchesi	Offline Operations Status
Moore	Accelerator Status and Plans
Nurse	Trigger and High Luminosity

http://cdorg.fnal.gov/rex/status%20report/20070416/20070416.htm

Data size and Storage assumptions

		size		tape factor	tier disk factor
sizes	raw event size	0.3	MB	1	0.001
	raw/reprocessing siz	0.5	MB	0.2	0.001
	data DST size	0.125	MB	1.2	0.1
	data TMB size	0.0125	MB	3	1
	data rootuple size	0.01	MB	0	0
	MC D0Gstar size	0.7	MB	0.1	0
	MC D0Sim	0.3	MB	0	0
	MC DST size	0.2	MB	0	0
	MC TMB size	0.02	MB	2	0.5
	PMCS MC size	0.02	MB	2	0.5
	MC rootuple size	0.02	MB	0	0

data samples (events)

	1 day	1 year	phase 1 2 years	phase 2 4 years
event rate	1.90E+06	6.94E+08	1.39E+09	8.33E+09

TAPE data accumulation (TB)

raw event	0.57	208.14	416.28	2497.65
raw/reprocessing	0.19	69.38	138.76	832.55
data DST	0.29	104.07	208.14	1248.83
data TMB	0.07	26.02	52.03	312.21
data rootuple	0.00	0.00	0.00	0.00
MC D0Gstar	0.13	48.57	97.13	582.79
MC D0Sim	0.00	0.00	0.00	0.00
MC DST	0.00	0.00	0.00	0.00
MC TMB	0.08	27.75	55.50	333.02
PMCS MC	0.08	27.75	55.50	333.02
MC rootuple	0.00	0.00	0.00	0.00
total storage (TB)	**1**	**512**	**1,023**	**6,140**
total storage (PB)	**0.001**	**0.51**	**1.02**	**6.14**
total storage (GB)	**1,402**	**511,672**	**1,023,343**	**6,140,059**

TIER DISK data accumulation (TB)

raw event	0.00	0.21	0.42	2.50
raw/reprocessing	0.00	0.35	0.69	4.16
data DST	0.02	8.67	17.34	104.07
data TMB	0.02	8.67	17.34	104.07
data rootuple	0.00	0.00	0.00	0.00
MC D0Gstar	0.00	0.00	0.00	0.00
MC D0Sim	0.00	0.00	0.00	0.00
MC DST	0.00	0.00	0.00	0.00
MC TMB	0.02	6.94	13.88	83.26
PMCS MC	0.02	6.94	13.88	83.26
MC rootuple	0.00	0.00	0.00	0.00
total storage (TB)	**0**	**32**	**64**	**381**
total storage (PB)	**0.000**	**0.03**	**0.06**	**0.38**
total storage (GB)	**87**	**31,776**	**63,551**	**381,308**

Farm Processing

Average Rate:	75
Farm Efficiency:	70%
Misc. Processing:	10%
Reprocessing:	50%
Cost/node:	3,000
I/O Cost/100 nodes	25,000

CPU	SpecI2000
3GHz	960
4GHz	1280
6GHz	1920
9GHz	2880
14GHz	4480
20GHz	6400

FY05 Target Spending Fraction:		20%		30%		50%		Total	
Execution	500MHz CPUs at	FY03, 3GHz Nodes		FY04, 4GHz Nodes		FY05, 6GHz Nodes		Target	
Time	Beginning of Run	No. Nodes	Cost	No. Nodes	Cost	No. Nodes	Cost	No. Nodes	Cost
30	5143	72	241,000	108	349,000	180	590,000	360	1,180,000
75	12857	180	590,000	271	888,000	452	1,481,000	903	2,959,000
100	17143	241	798,000	361	1,183,000	602	1,981,000	1204	3,962,000

Farm processing capacity in Summer ’02 ~50Hz

D0mino backend
16 node, 1 GHz
80 nodes, 2 GHz, summer ‘02

Summary of infrastructure costs:

Infrastructure Costs	2003	2004	2005	2006	2007	Total
Databases:						
DB Hosts, Sun, then Linux	$60,000	$60,000	$25,000	$25,000	$25,000	
non COTS disk and controllers	$60,000	$20,000	$10,000	$10,000	$10,000	
Mirrors	$30,000	$15,000	$25,000	$15,000	$15,000	
Software	$50,000	$0	$50,000	$0	$50,000	
DB totals	$200,000	$95,000	$110,000	$50,000	$100,000	$555,000
Networking	$80,000	$50,000	$100,000	$200,000	$400,000	$830,000
Build Machines	$50,000	$50,000	$50,000	$50,000	$50,000	$250,000
Additional SAM servers	$50,000	$50,000	$50,000	$50,000	$50,000	$250,000
Total, fixed cost	$380,000	$245,000	$310,000	$350,000	$600,000	$1,885,000

Rate assumptions

rates	average event rate	22	Hz
	raw data rate	22.5	MB/s
	Geant MC rate	11	Hz

Average rate assumes an accelerator and experiment Duty factor applied to a peak rate of 50 Hz

rate increase assumptions		
	rate factor	3
	phase_1	2
	phase_2	4
	last year	2009
	total years	6
	down year	2005

Full Cost Estimate, No I/O replacement

Extremely preliminary D0 C&S cost estimate								
		2002	2003	2004	2005	2006	2007	***Total(2003-2007)***
Fixed Infrastructure Costs		$400,000	$380,000	$245,000	$310,000	$350,000	$600,000	$1,885,000
farm + analysis cpu		$800,000	$640,000	$938,000	$1,531,000	$500,000	$500,000	$4,109,000
disk cache		$0	$150,000	$100,000	$50,000	$150,000	$150,000	$600,000
robotic storage		$400,000	$150,000	$0	$150,000	$150,000	$150,000	$600,000
tape drives		$200,000	$600,000	$300,000	$300,000	$600,000	$600,000	$2,400,000
D0mino upgrade		$150,000	$0				$0	$0
Backup facility			$350,000					
Sum		$1,950,000	$2,270,000	$1,583,000	$2,341,000	$1,750,000	$2,000,000	$9,944,000

Update on D0 and CDF computing models and experience

Amber Boehnlein
FNAL/CD
For CDF and DO collaborations
June 27, 2003

Vital Statistics

Vital Statistics	CDF	DO
Raw Data Size (kbytes/event)	205	230(160)
Reconstructed Data Size (kbytes/event)	180	200
Primary User data (kbytes/event)	N/A	20
User Skims	DST	TMB
User Skims(kbytes/event)	25-180	20-40
Reconstruction Time (Gh-sec/event)	4	20
Monte Carlo Chain	fast	full Geant
Peak Data Rate(Hz)	75	50
Persistent format	RootIO	D0om/dspack

D0 Vital Statistics 2006

D0 Vital Statistics		
	1997(projections)	2006
Peak (Average) Data Rate(Hz)	50(20)	100(35)
Events Collected	600M/year	1.5 B
Raw Data Size (kbytes/event)	250	250
Reconstructed Data Size (kbytes/event)	100 (5)	80
User format (kbytes/event)	1	40
Tape storage	280 TB/year	1.6 pb on tape
Tape Reads/writes (weekly)		30TB/7TB
Analysis/cache disk	7TB/year	220 TB
Reconstruction Time (Ghz-sec/event)	2.00	50 (120)
Monte Carlo Chain	full Geant	full Geant
user analysis times (Ghz-sec/event)	?	1
user analysis weekly reads	?	3B events
Primary Reconstruction farm size (THz)	0.6	2.4 THz
Central Analysis farm size (GHz)	0.6	2.2 THz
Remote resources(GHz)	?	~ 2.5 THz(grid)

DO Analysis-2003

D0 Analysis systems

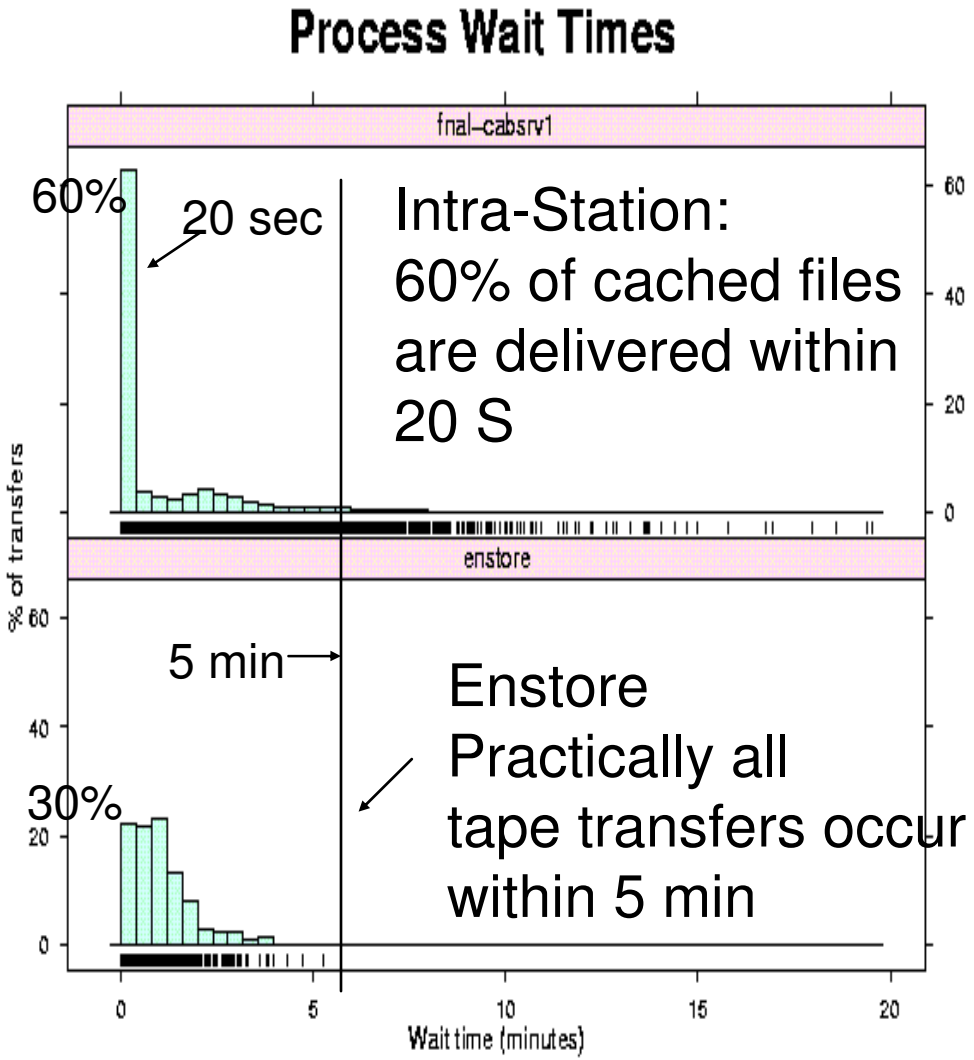

User interface including batch submission –D0tools

CLUED0-managed by the users for the users

- **Clustered desktops with batch system and SAM station, local project disk**
- **Developed expertise and knowledge base**

Linux fileservers and worker nodes for analysis

- **pioneered by CDF with FNAL/CD**

Before adding 100 TB of Cache, 2/3 of transfers could be from tape.

Central Robotics

20TB
At peak

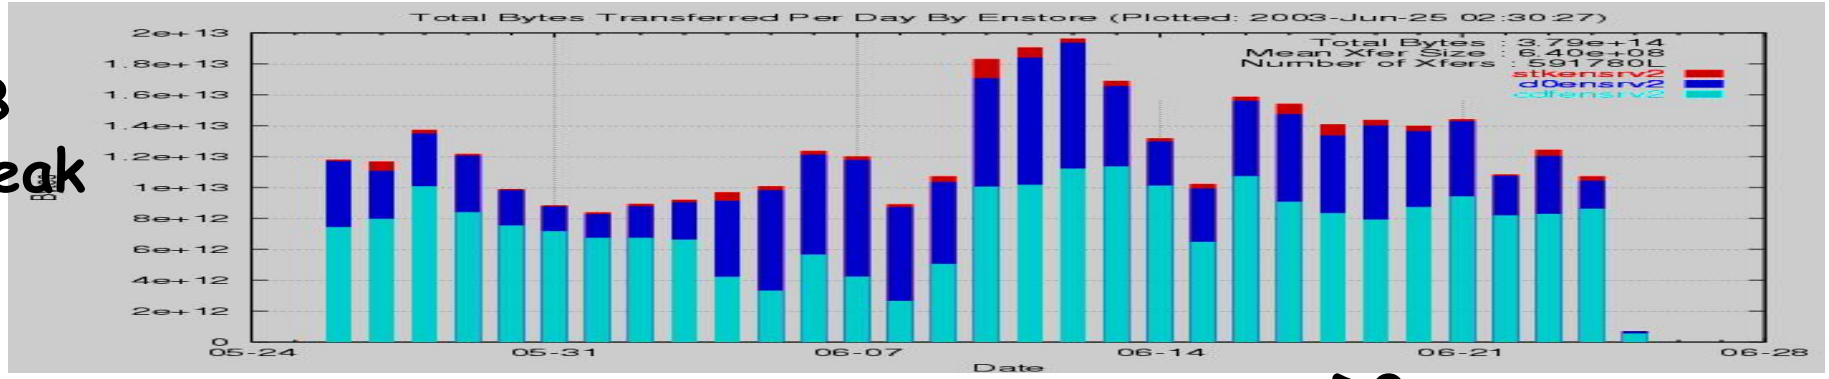

CDF Data to tape, June 25, 2003 D0

Library	Stored	#tapes
9940a	302TB	5521
9940b	104TB	1046

Library	Stored	#tapes
STK	219TB	3780
9940b	30.7TB	380
LTO	87.6TB	1099

CDF Drive usage

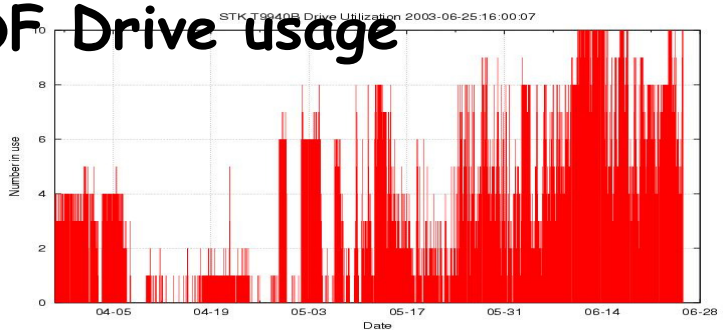

Known data loss due to Robotics/ Enstore for DO—3 GB—Seriously.

2004 D0 Production

Offsite Monte Carlo generation from 1999 !

In 2004 1M event/week at 6 sites

The reconstruction speed is a bottleneck.

Offsite reconstruction from raw data was clearly going to be essential

FNAL Farm scripts not transportable

Develop centralized job submission for execution on the Grid
(Job and Information Monitoring (JIM)) <SAMGrid>

2.5M

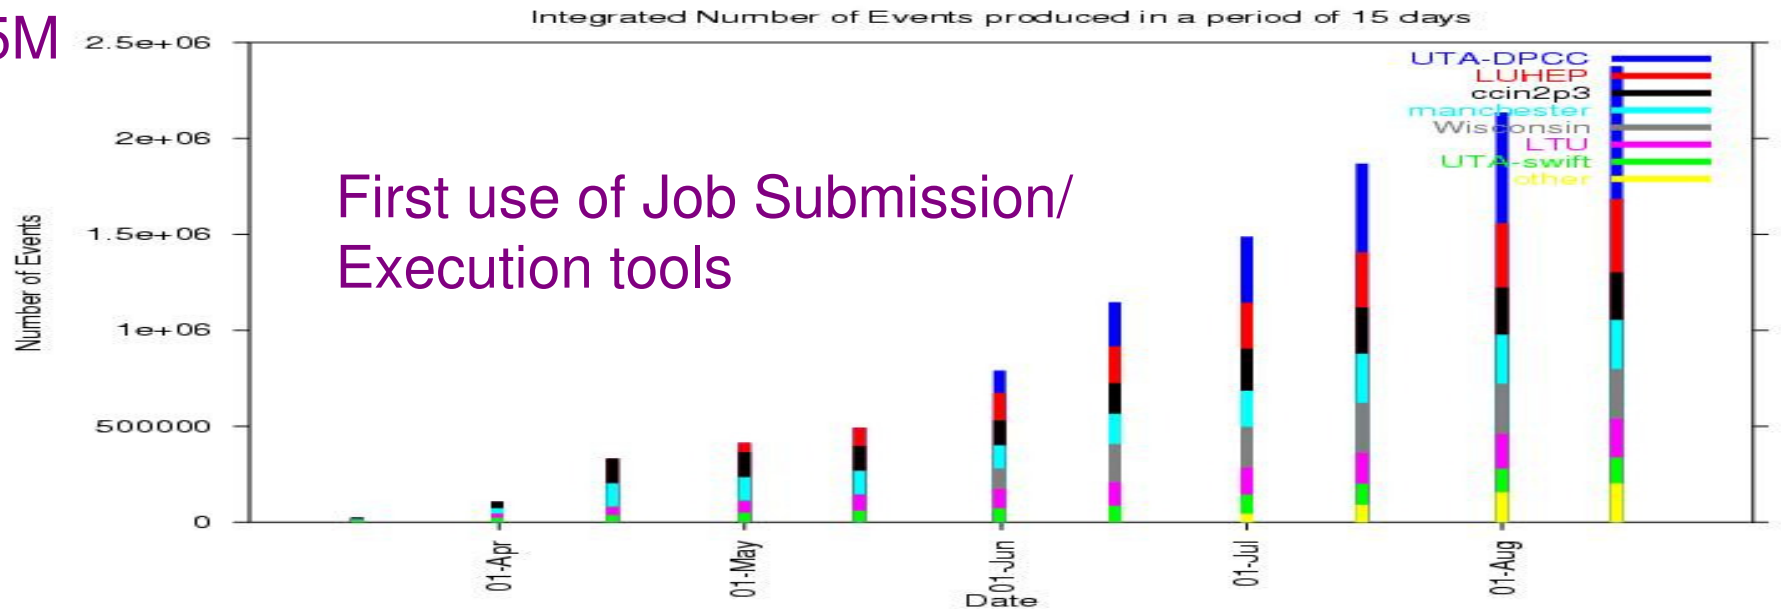

Mar - Nov 2005 D0 Reprocessing

Six months development and preparation

1B events from raw – SAMGrid default – basically all off-site

Massive task largest HEP activity on the grid

~3500 1GHz equivalents for 6 months

200 TB

Largely used shared resources – LCG (and OSG)

P17 Reprocessing Status as of 24-Nov-2005 (all sites)

Total Raw Events	986190444	
Processed Events	958741259	
Sites	fnal FNAL OSCER FZU_GRID WestGrid ccin2p3 GridKa UTA-DPCC Wisconsin IMPERIAL_PRD CMS-FNAL-WC1 SPRACE	

P17 Reprocessing Status as of 24-Nov-2005 (Remote sites only)

Processed Events	821900405	
Sites	fnal FNAL OSCER FZU_GRID WestGrid ccin2p3 GridKa UTA-DPCC Wisconsin IMPERIAL_PRD CMS-FNAL-WC1 SPRACE	

2006 D0 Monte Carlo Production

Significantly increased production as ramp up SAMGrid usage

- 150M events in last 6 months
- Up to 10M/week
- 20M on LCG via Nikhef
- Full LCG interoperability now being commissioned

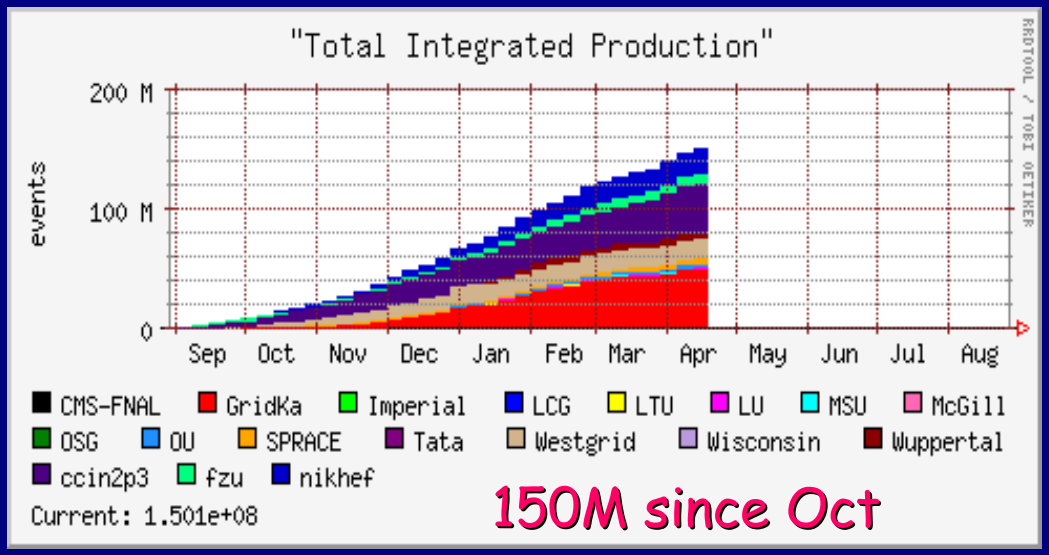

Up to 10M / week

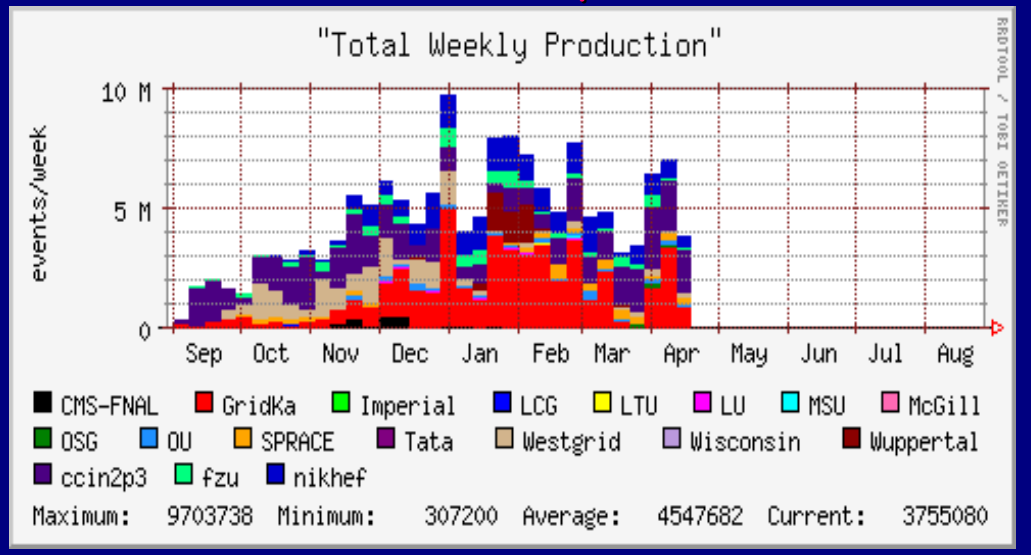

2006 Production "Fixing"

December 2005-problem found with a hadronic calorimeter correction

- Mobilized to "fix" 1.4 Billion events in six weeks
- Increased interoperability
  - Extensive use of OSG (CMS farm) (OSG with local SAMGrid installation)
  - Full LCG interoperability (without local SAMGrid)
  - First use of OSG facilities without a local SAMGrid installation
- Improved data quality monitoring
- Finished early—Great advertisement for the Grid

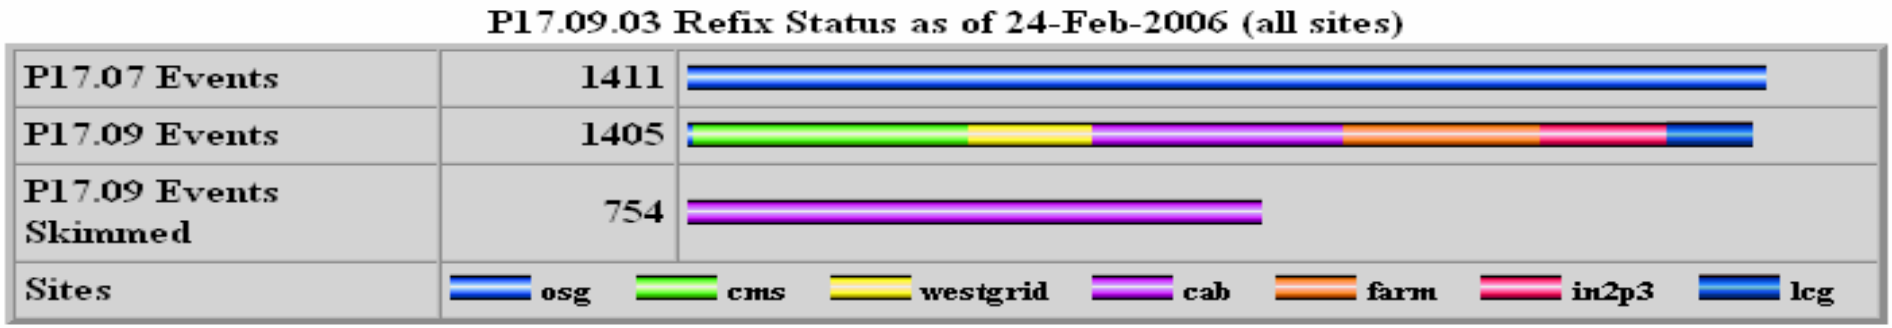

2006 D0 Analysis

Evolution in data tiers

In 2004, common root data tier- Common Analysis Format "CAF" Project begins

CAF commissioned in 2006 use taking off

Working to understand use cases, Next focus is analysis

User platform access (in M events)
Red is TMB access
Blue is CAF
Black is Physics group samples

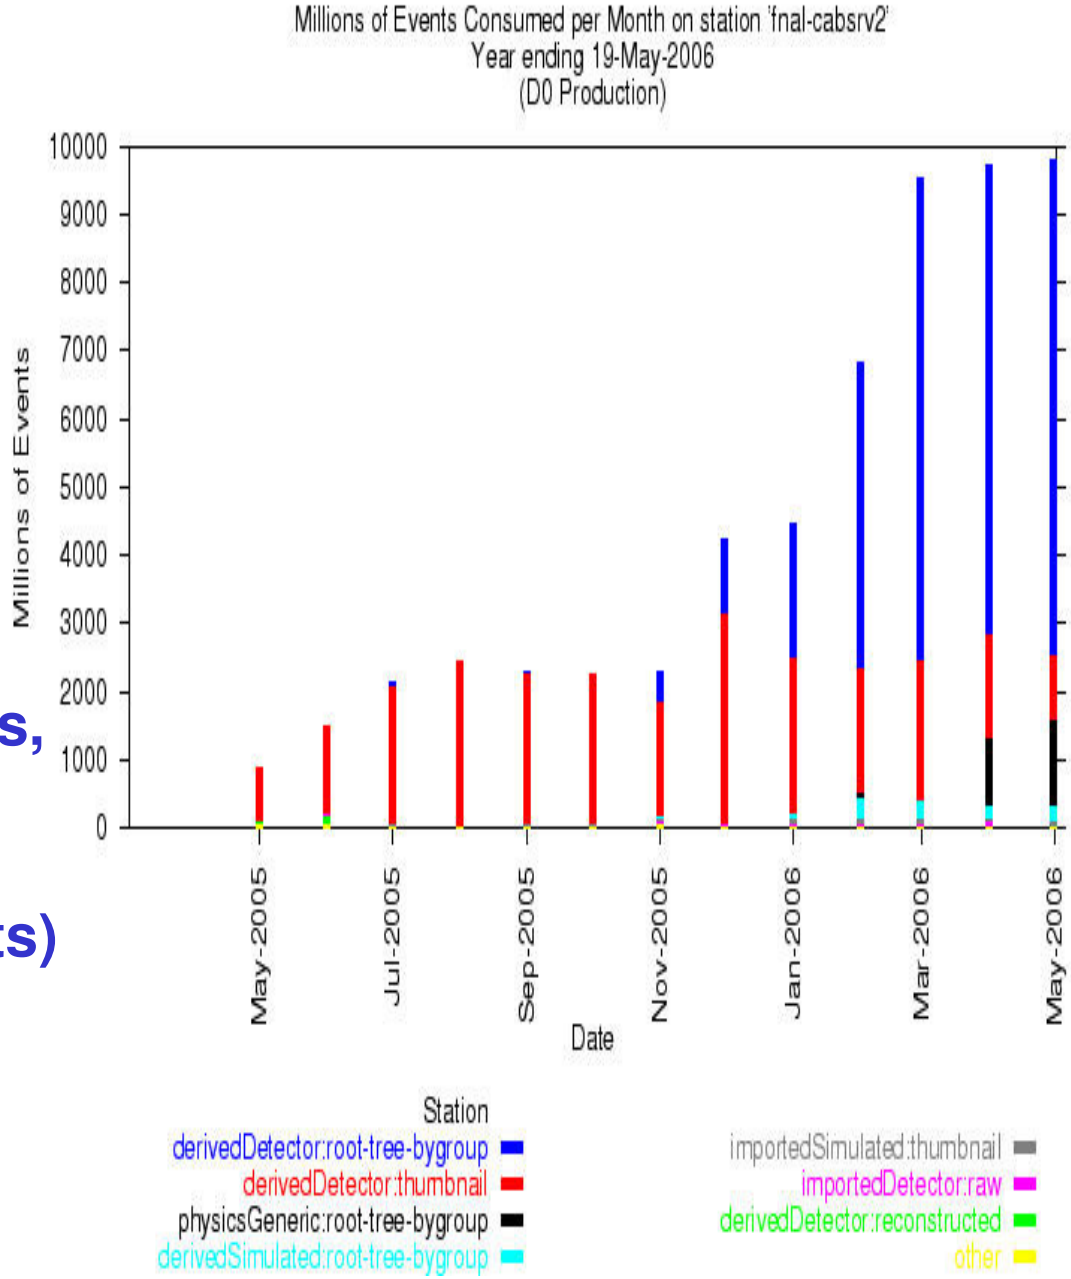

D0 Analysis 2006

Events consumed by station since "the beginning of SAM time"

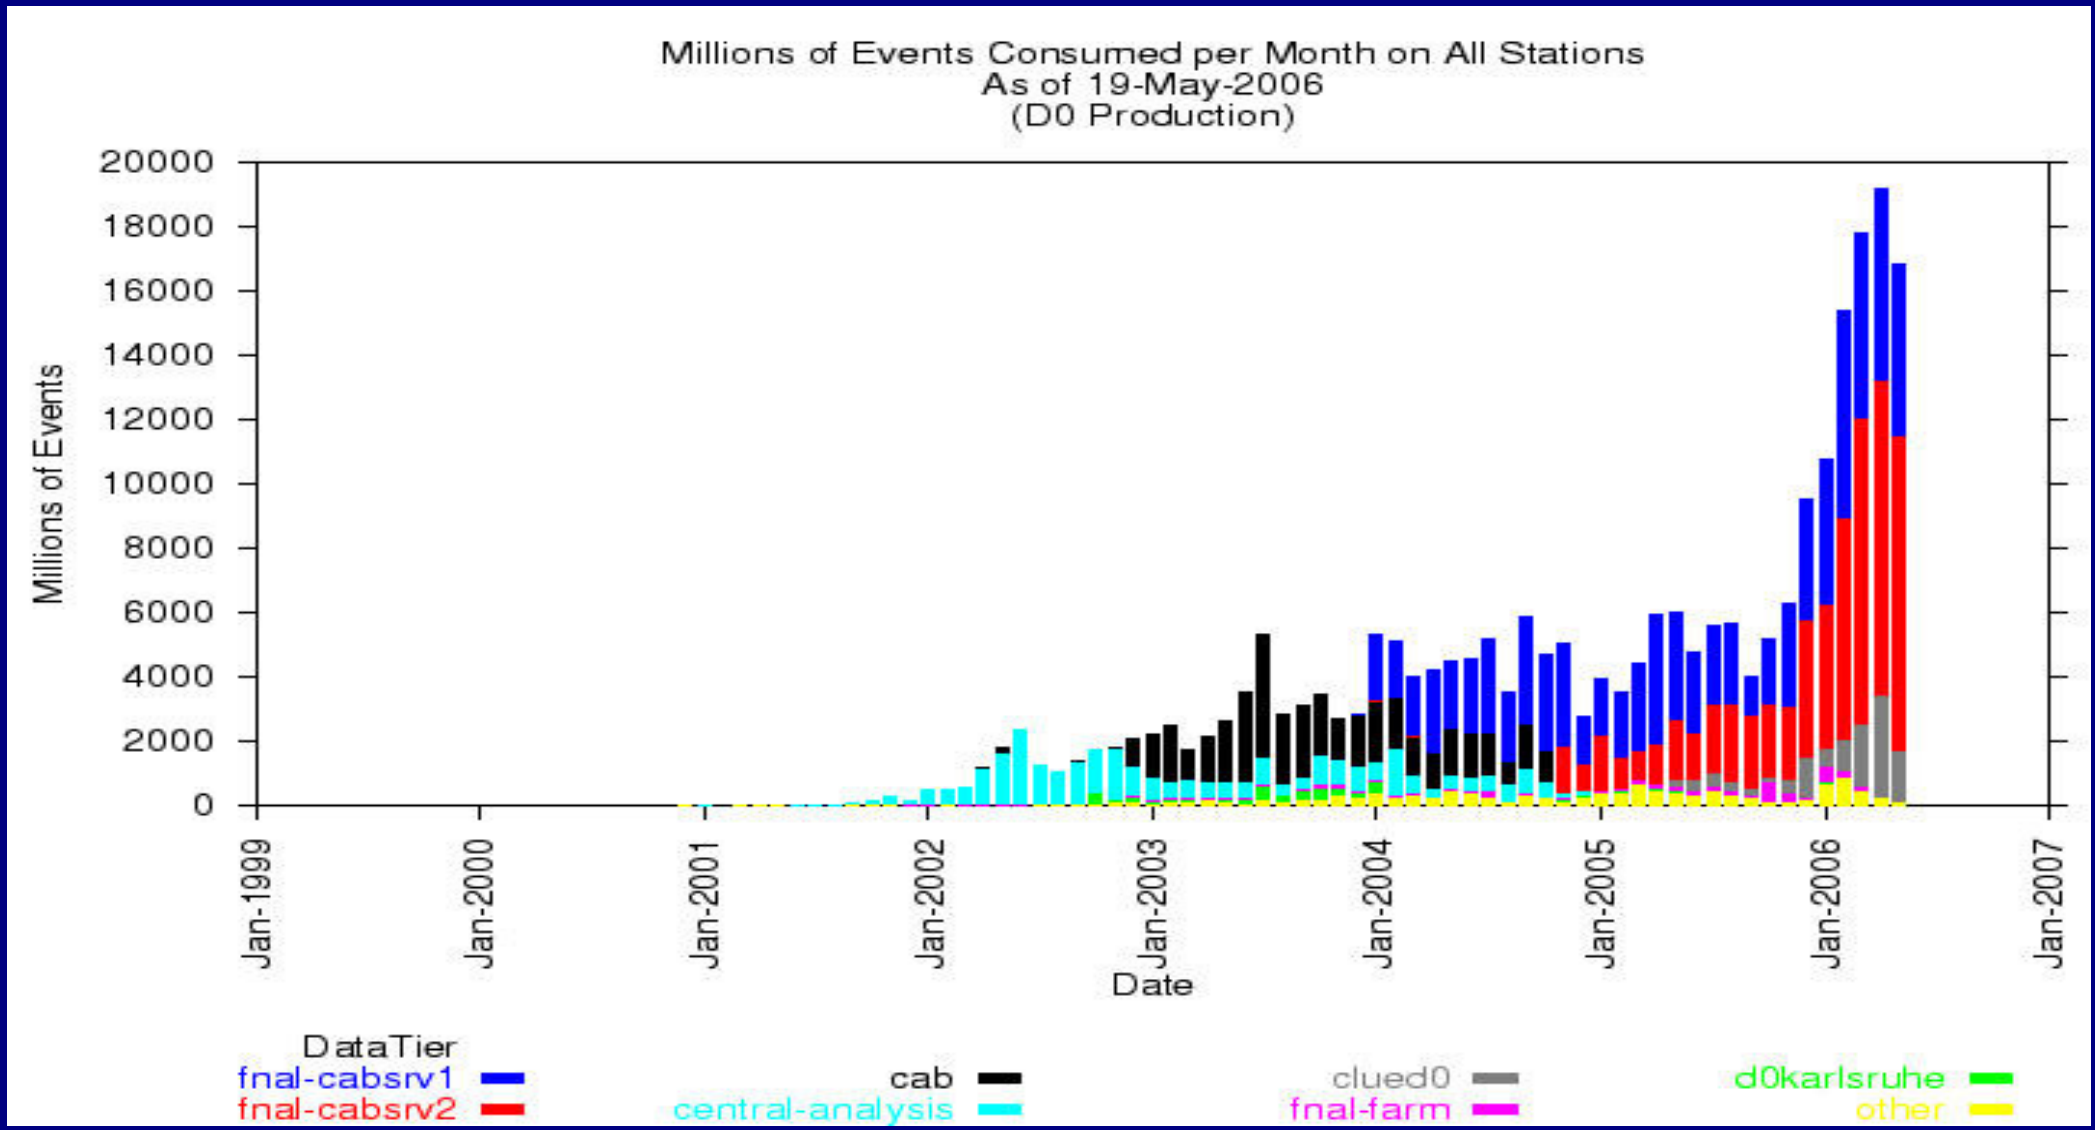

CDF Publication History

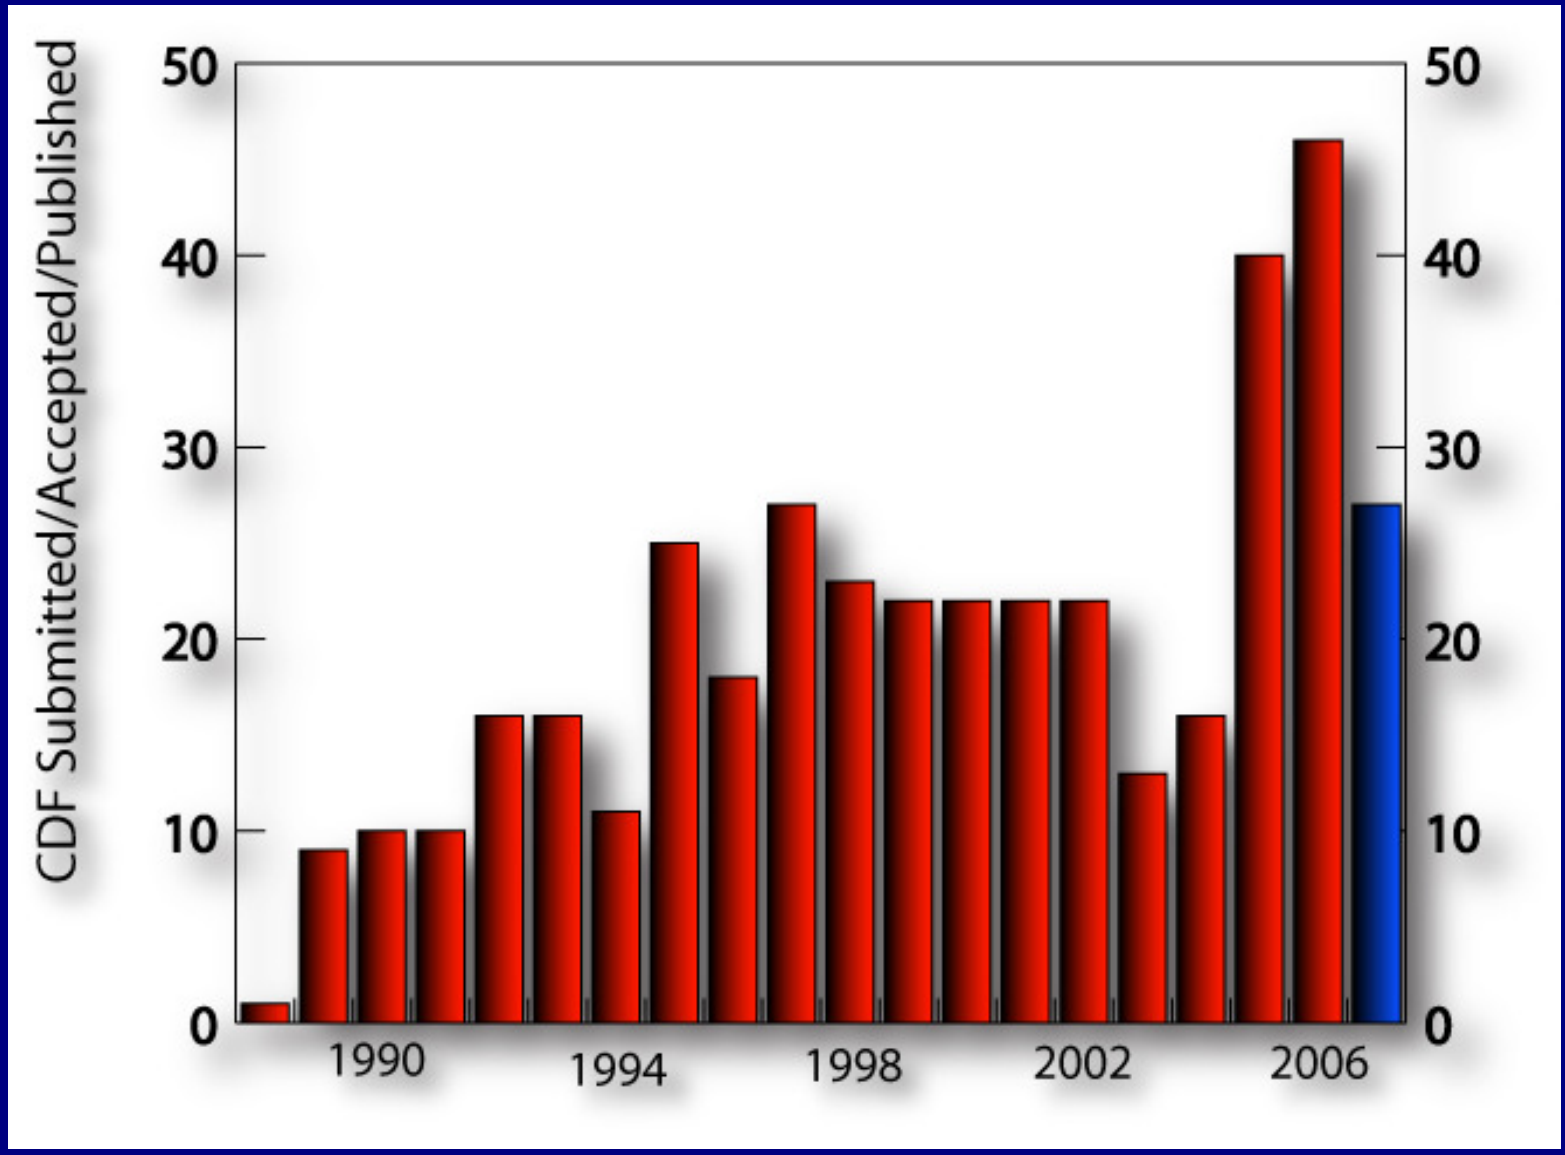

Computing at CDF: setting the scale

- Data rate from the detector: 200 Hz increasing to over 300 Hz, corresponding to 40-60 MB/s
- reconstructed: ~ 5 x 10^9 events, corresponding to a data volume > 1.5 PetaBytes
- fully simulated > 10^9 events
- ~ 600 physicists doing analysis around the world

Data Reconstruction

- Event size and processing time depend on type of trigger and inst. luminosity
- Output from detector divided into 8 data "streams"
- Output of production divided in over 50 physics datasets
- Typical time to reconstruct one event: ~3 seconds (mostly tracking)
- Event rate ~200 Hz
- Need ~200 dual processors, ~1 THz
- Typical event size 150 kB

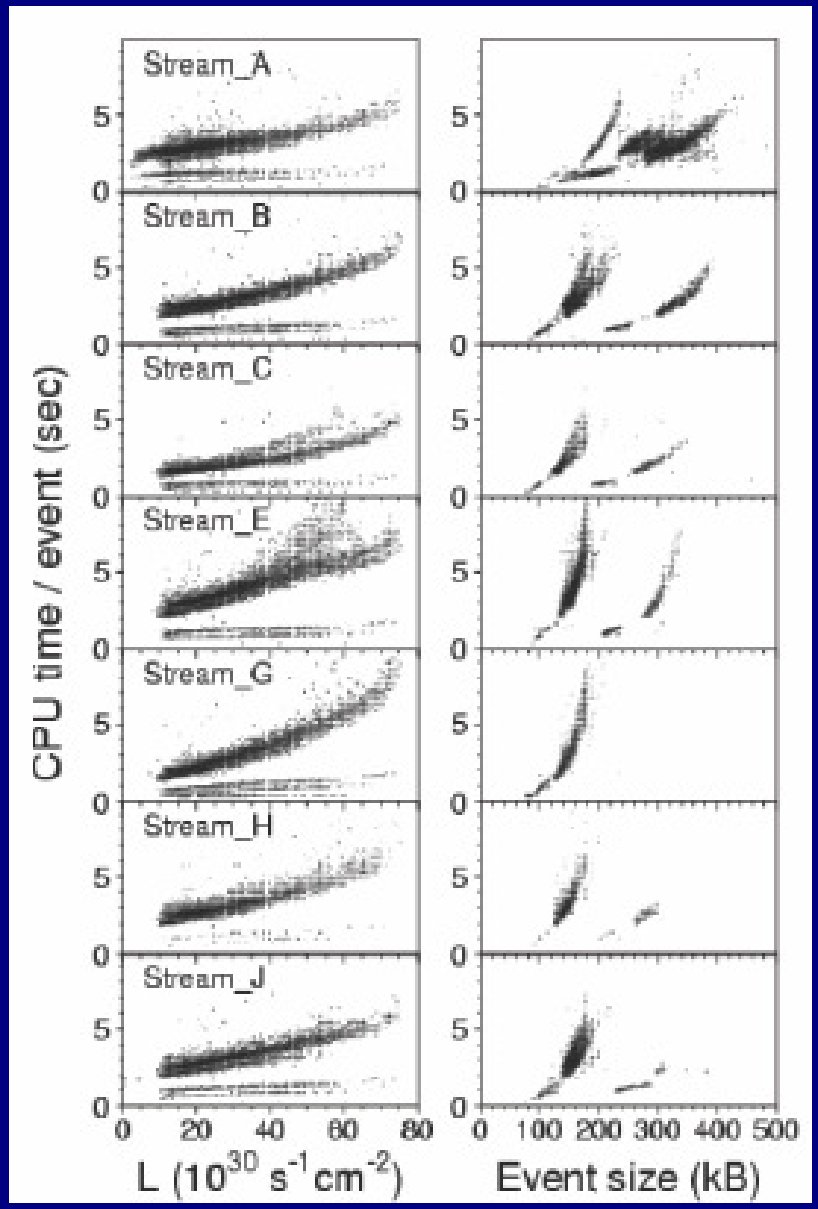

Data Processing Goals

- Pre-process part of the data to determine detector calibrations, alignment, beam positions
- Equivalent to processing 1.3 times the data
- Have delivered fully calibrated datasets 6 weeks after data taking
- Have achieved 25M events per day
(we need ~ 5M to keep up with incoming data)

Production Farm and Ntuples

- Main data representation for most CDF physicists is in the form of Ntuples (2-3 main flavors)
  - Very large user-driven ntupling tasks put extra load on data handling and analysis systems:
  - Serving needs of physics users is the hardest part of offline project
- We now also use the Production Farm to produce official ntuples.

CDF Analysis Farms (CAF)

- Main analysis platform for the experiment
- Contains the bulk of the CDF computing capacity
  - Fermilab CAF ~ 6 THz (mainly analysis tasks)
  - Remote CAFs ~ 2.5 THz (bulk of simulation tasks)
- Usage much less predictable than data reconstruction on Production Farm. Main tasks performed:
  - Secondary, tertiary dataset production
  - Ntuple production and analysis
  - Simulation (at remote sites) very cpu-intensive
- CAFs are serving the users very well

Remote Analysis Farms

- Fraction of CPUs used at various remote analysis facilities

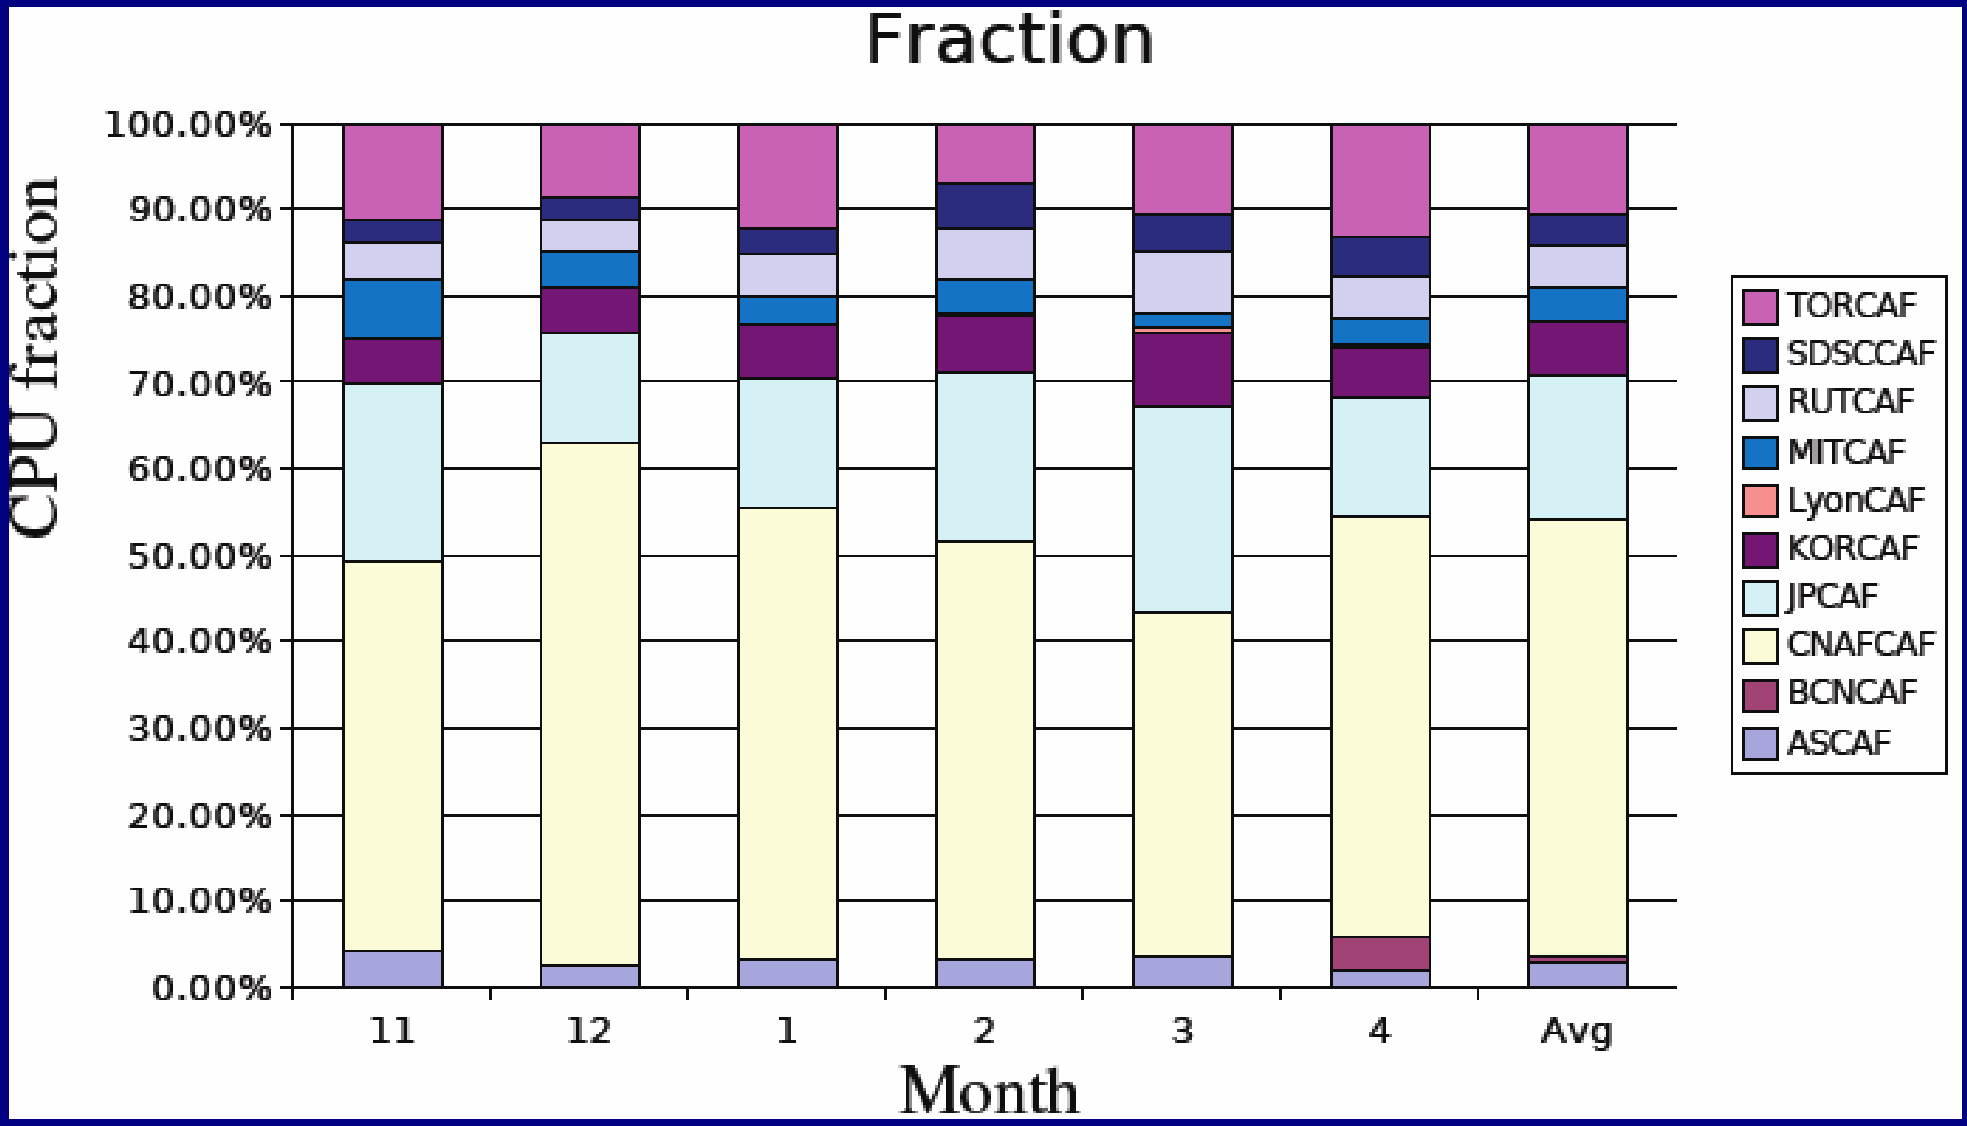

Analysis Farms Tasks

- Most CPUs used for analysis at FNAL

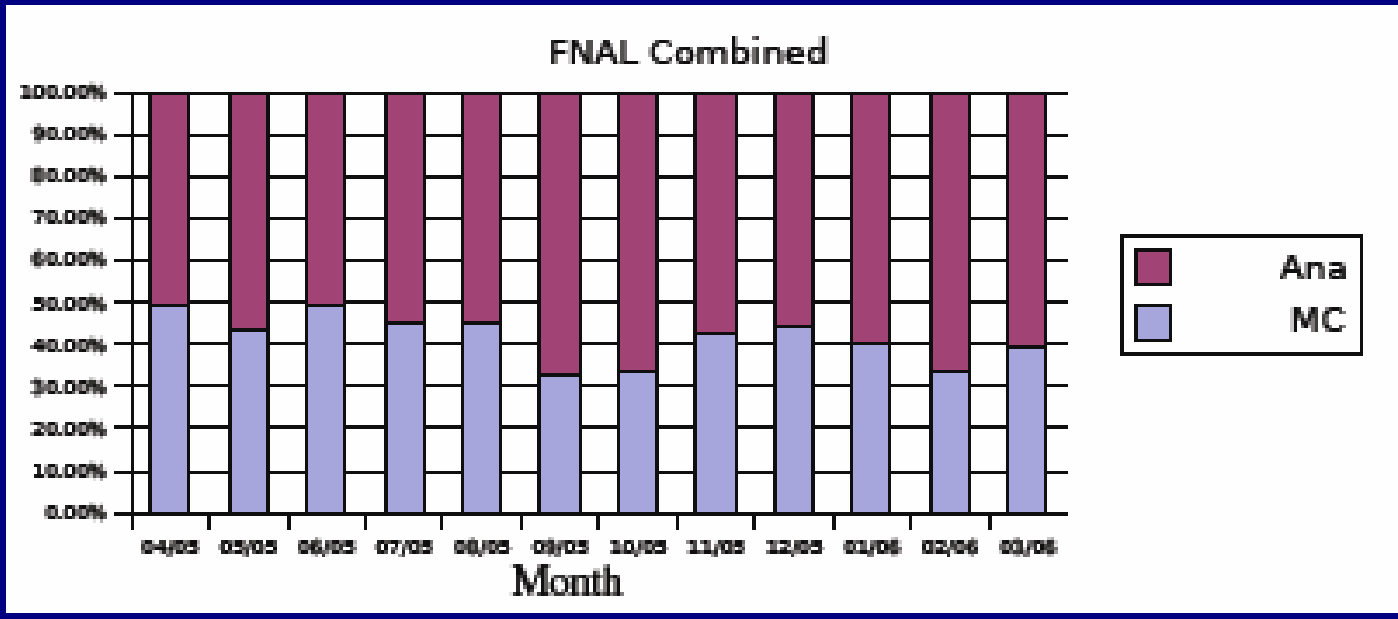

- Most CPUs used for simulation at remote facilities

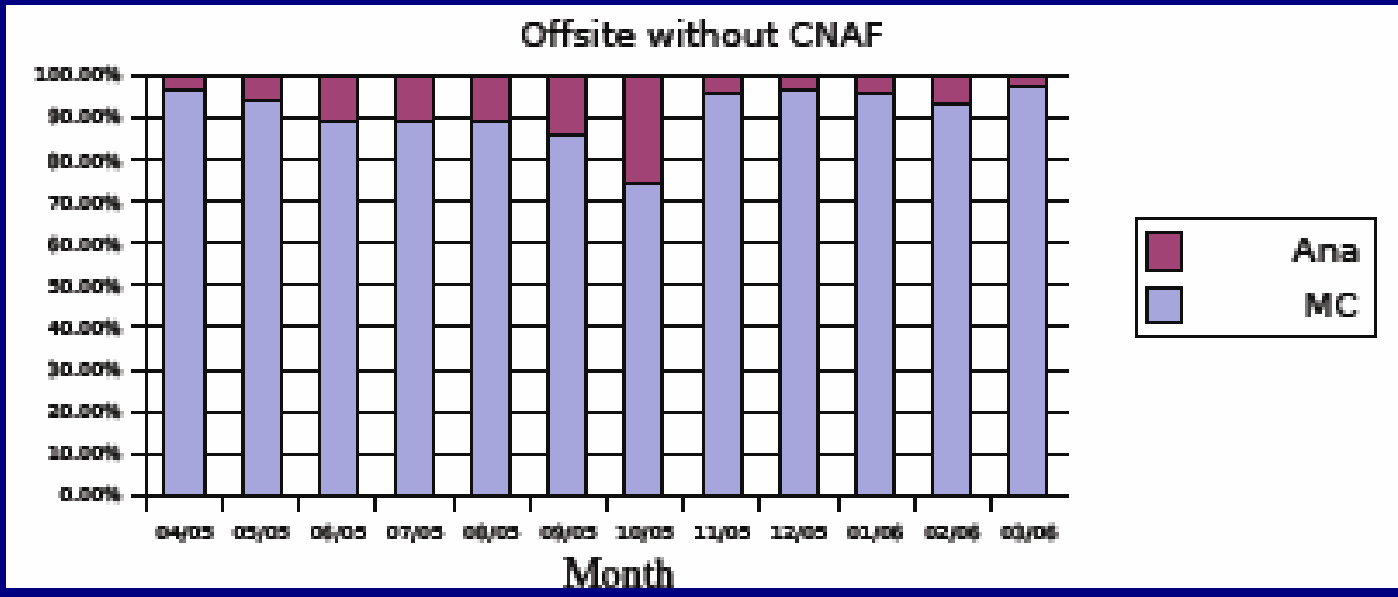

Analysis Tasks

- Time needed to unpack, read and do minimal analysis is < 0.05 sec/event
- User Analysis on production data average 0.75 sec/event
- 20% of tasks require more than 1 sec (40% of full reconstruction)
- Tail of distribution involves track refitting, vertex finding/fitting (dominated by B group)

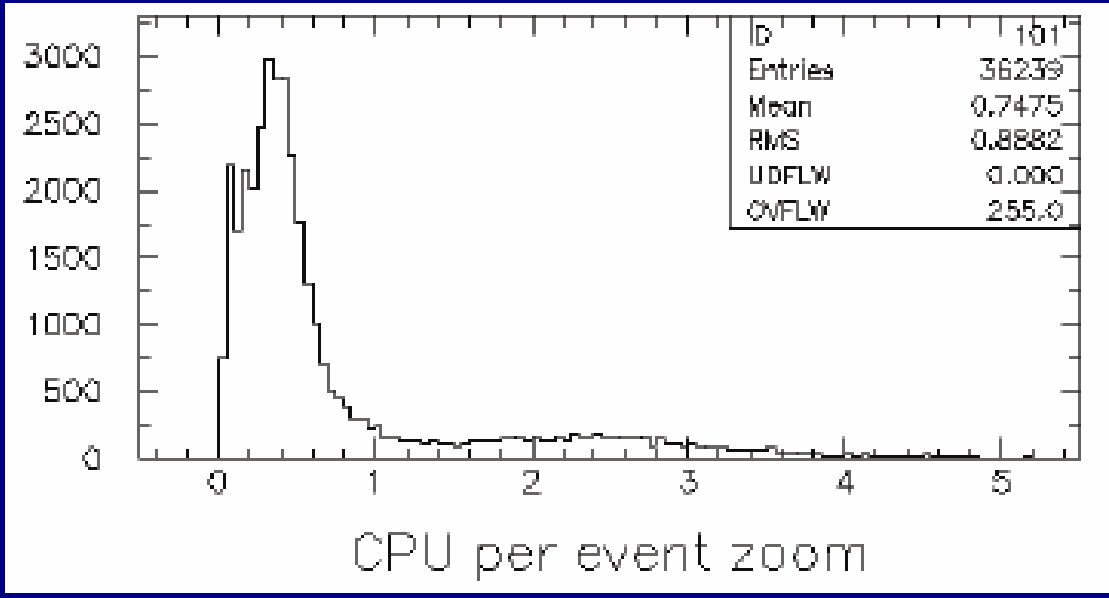

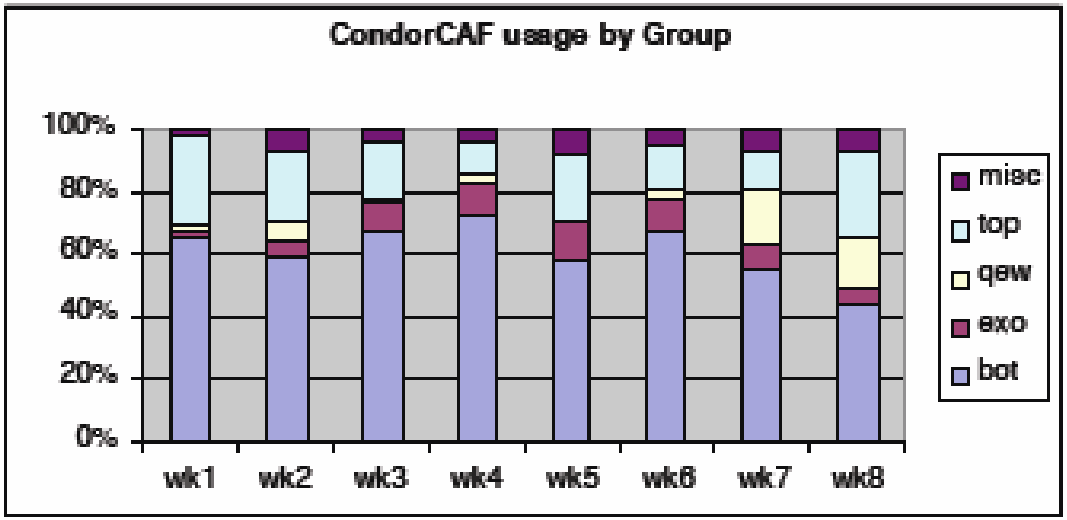

From 2004

Simulation

- For a very active event e.g. top: 7 secs on 2.5 GHz cpu
  - It would take > 10 minutes/event to do the same with full GEANT
- A clever mix of GEANT and parameterized simulation has enabled us to produce > 10^9 events
  - Dec04 to Aug05

    > 250M events produced in Canada, > 200M at Fermilab

Data Handling

Two main components: dCache and SAM

- dCache (joint project of Desy, FNAL):
  - "Virtualizes" disk used for local cache
    - Data on tape or distributed across local files servers
    - Exact location hidden from user
- Used only dCache and data catalog for more than 2 years

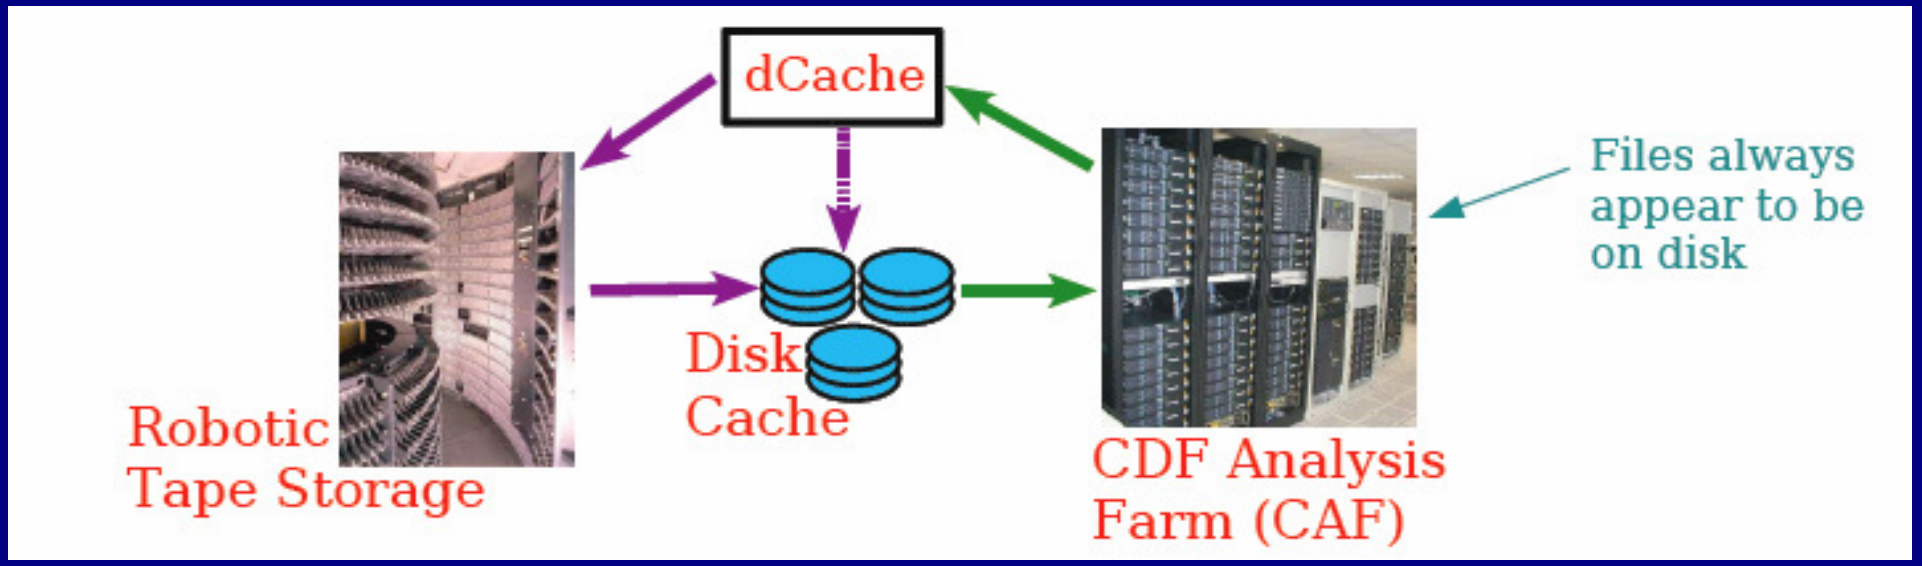

Data Handling (dCache)

- Data from dCache (average 10-25 TB a day)

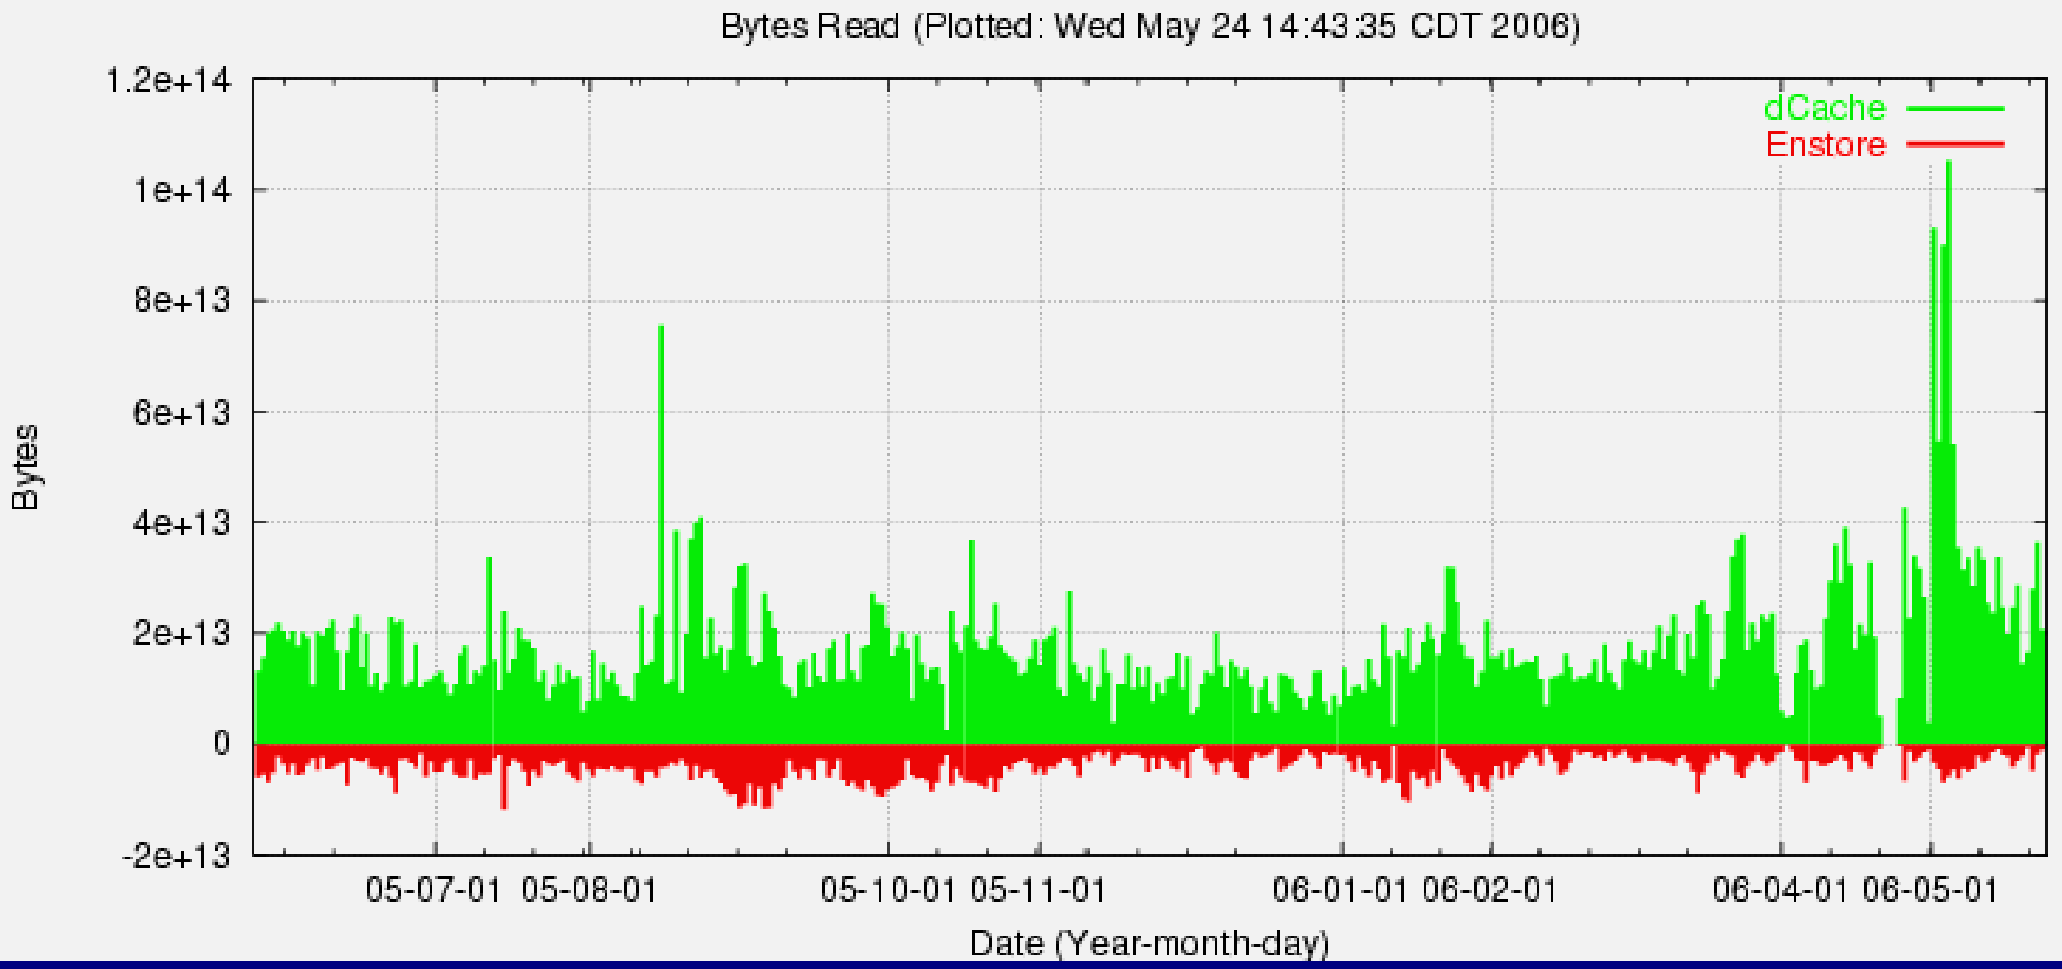

Data Handling (SAM)

- Files delivered by SAM at FERMILAB

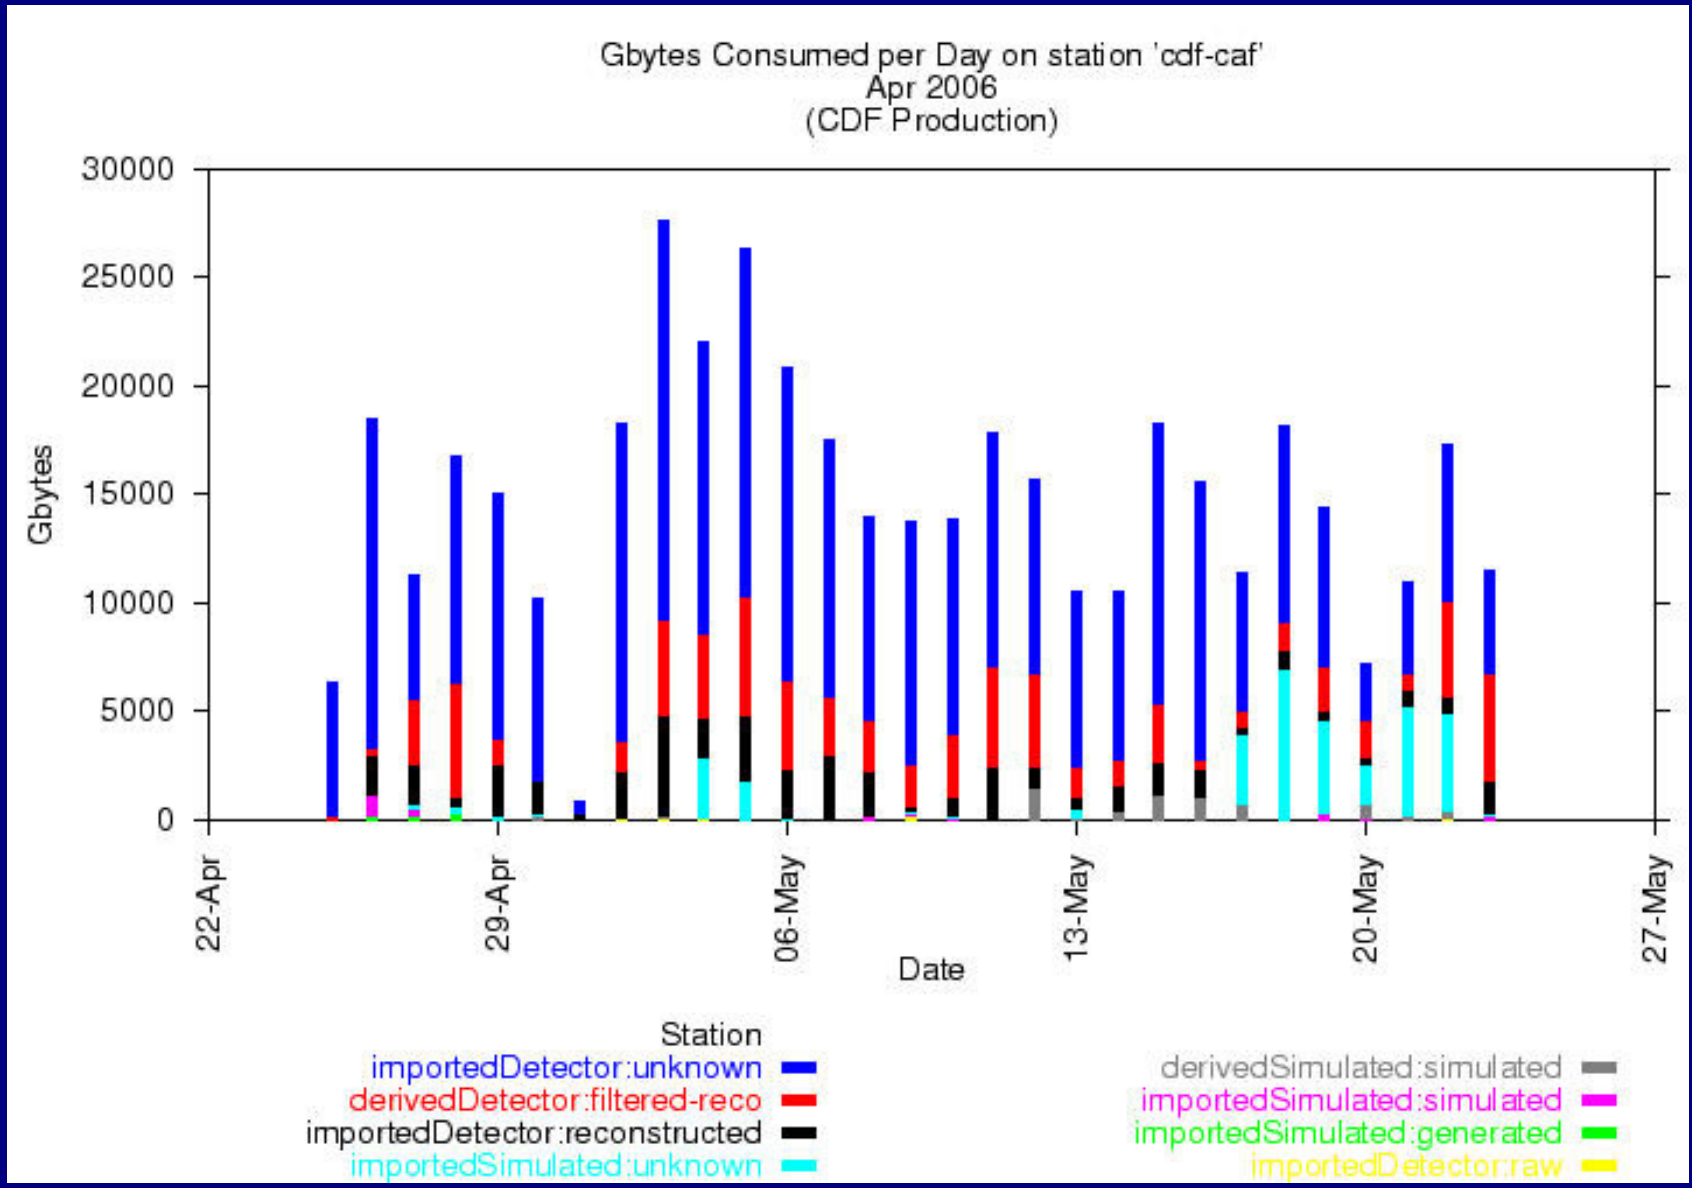

Future Challenges

- Higher instantaneous luminosity
  - Larger events, slower reconstruction, tracking more difficult, need more CPU per event
- Higher integrated luminosity and higher data taking rate
  - Larger data samples
    - Need more processing power
    - Need more storage
- Migration of physicists to LHC experiments
  - Human resources for operations are shrinking
- FY2010 Running has been proposed

Challenge I: Higher Inst.Luminosities

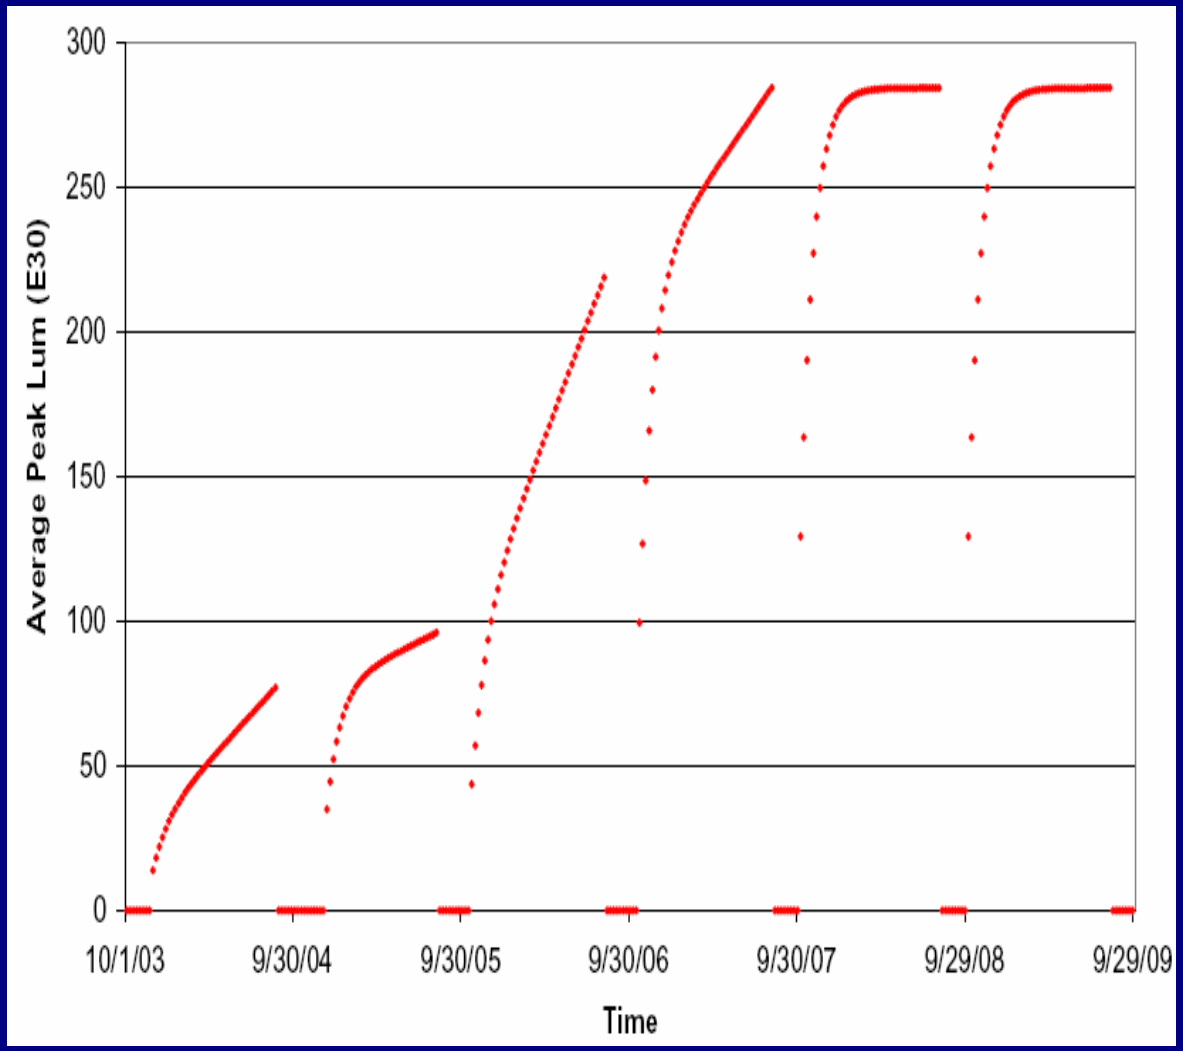

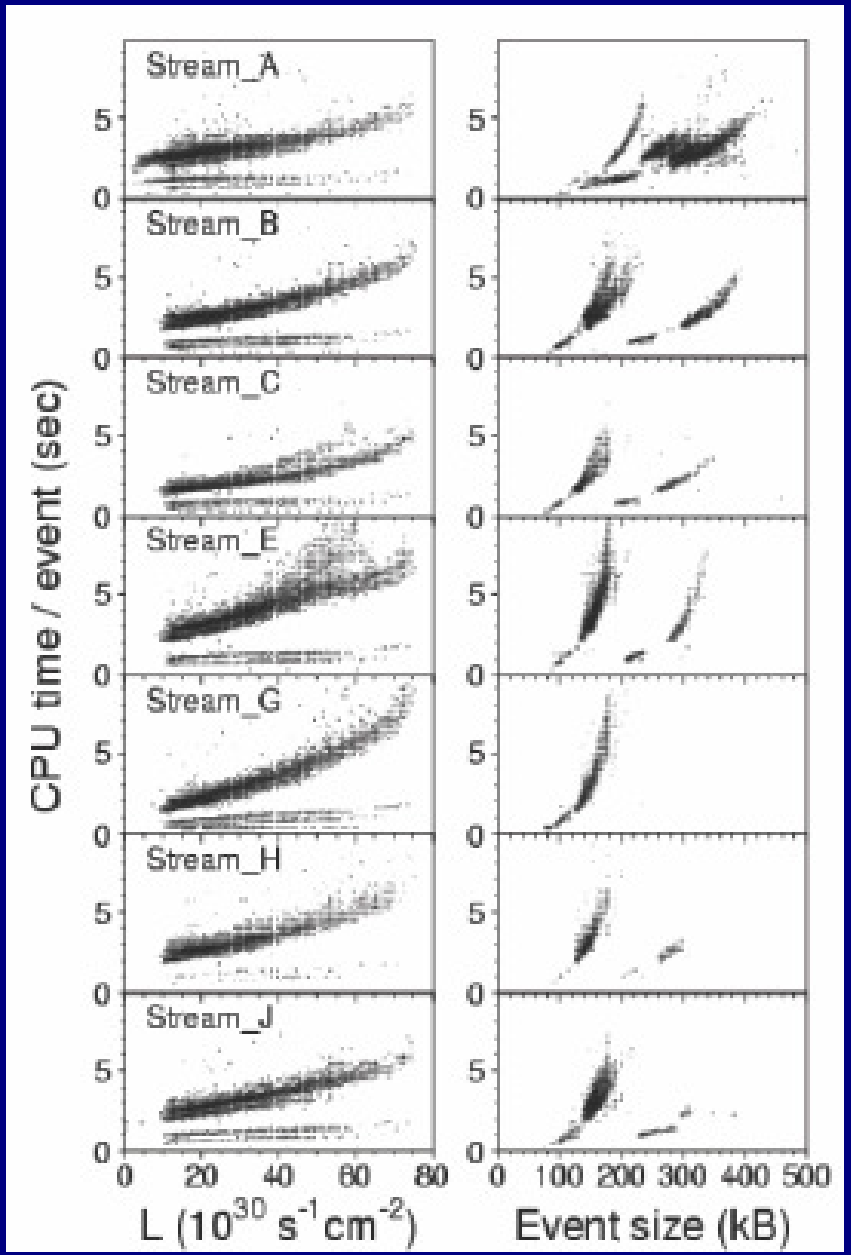

Challenge II: larger data samples

Fiscal Year	2006	2007	2008
Integrated luminosity (fb^{-1})	2.2	3.8	6.1
Total number of events (10^9)	3.4	5.7	9.2
Raw data logging rate (MB/s)	60*	60	60

* 40 MB/s until Q3

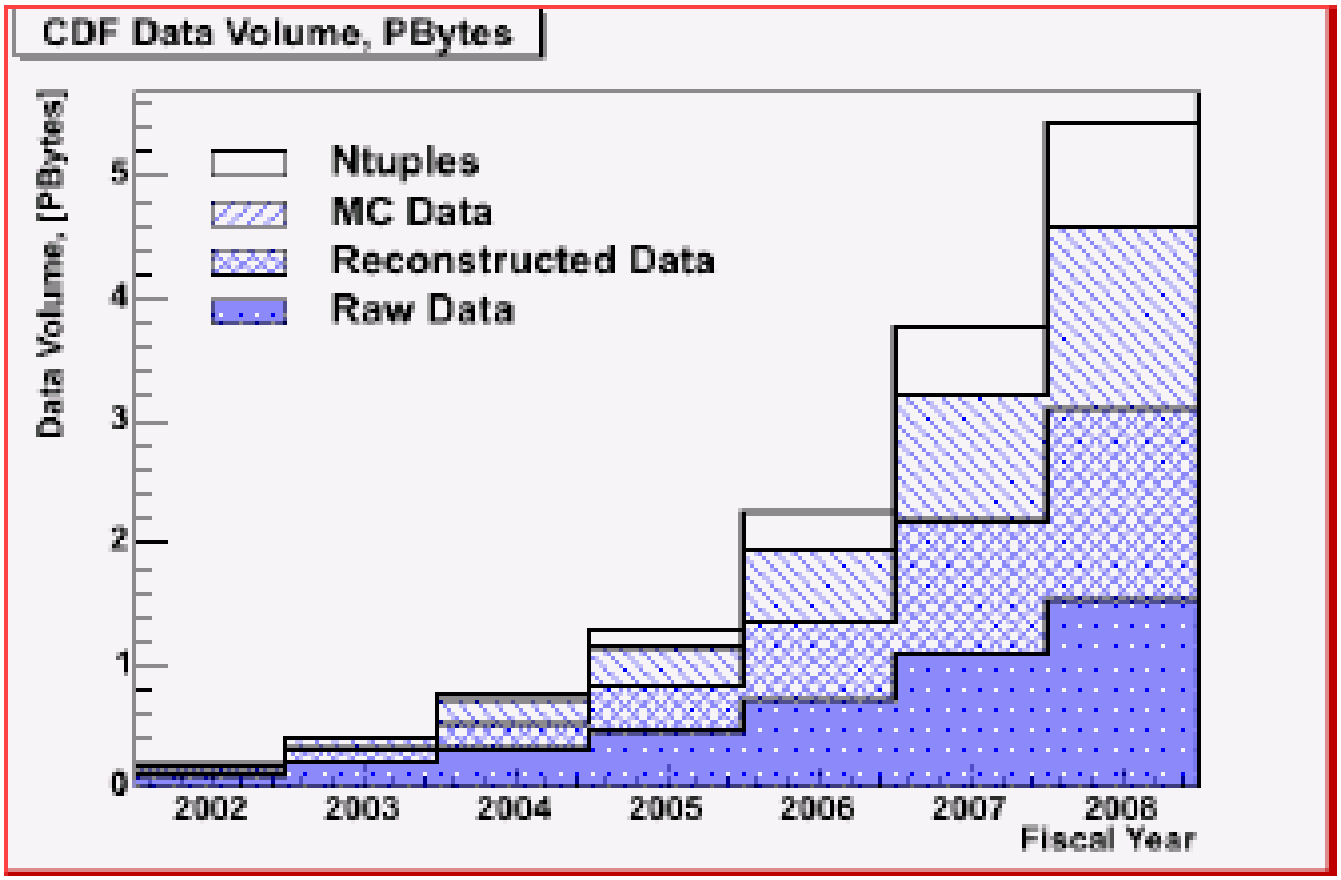

CDF

2003 2007

Total luminosity:

~2.9 fb^{-1} delivered, ~2.4 fb^{-1} to tape

Level 1 trigger: 12KHz ⇒ 35KHz

Level 2 trigger: 300Hz ⇒ 800Hz

Level 3 trigger: 24MB/s ⇒ 100MB/s

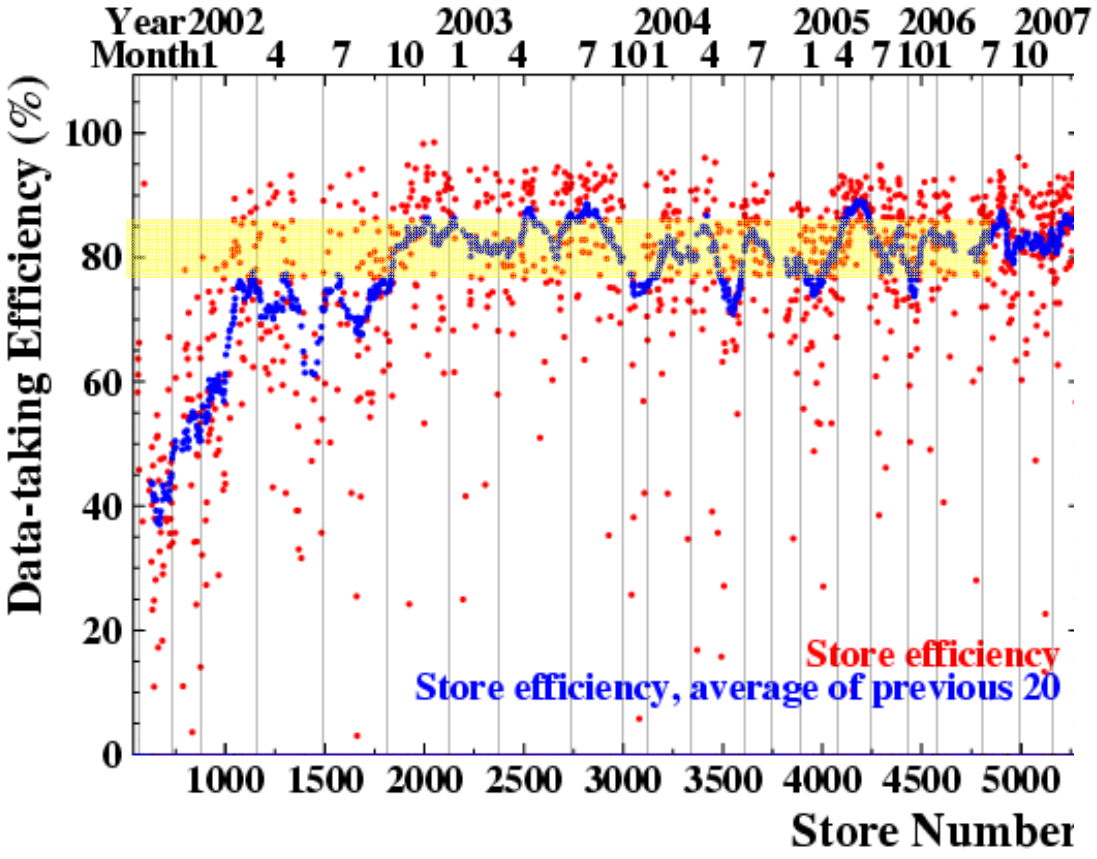

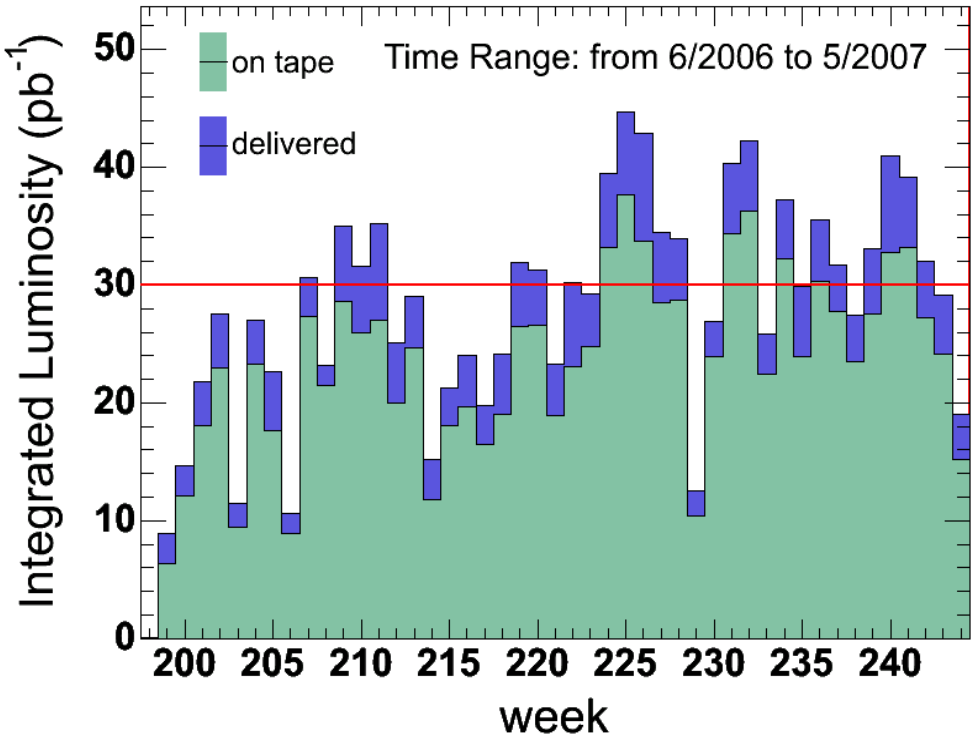

Dataset has doubled each of the last 4 years

Future Operation

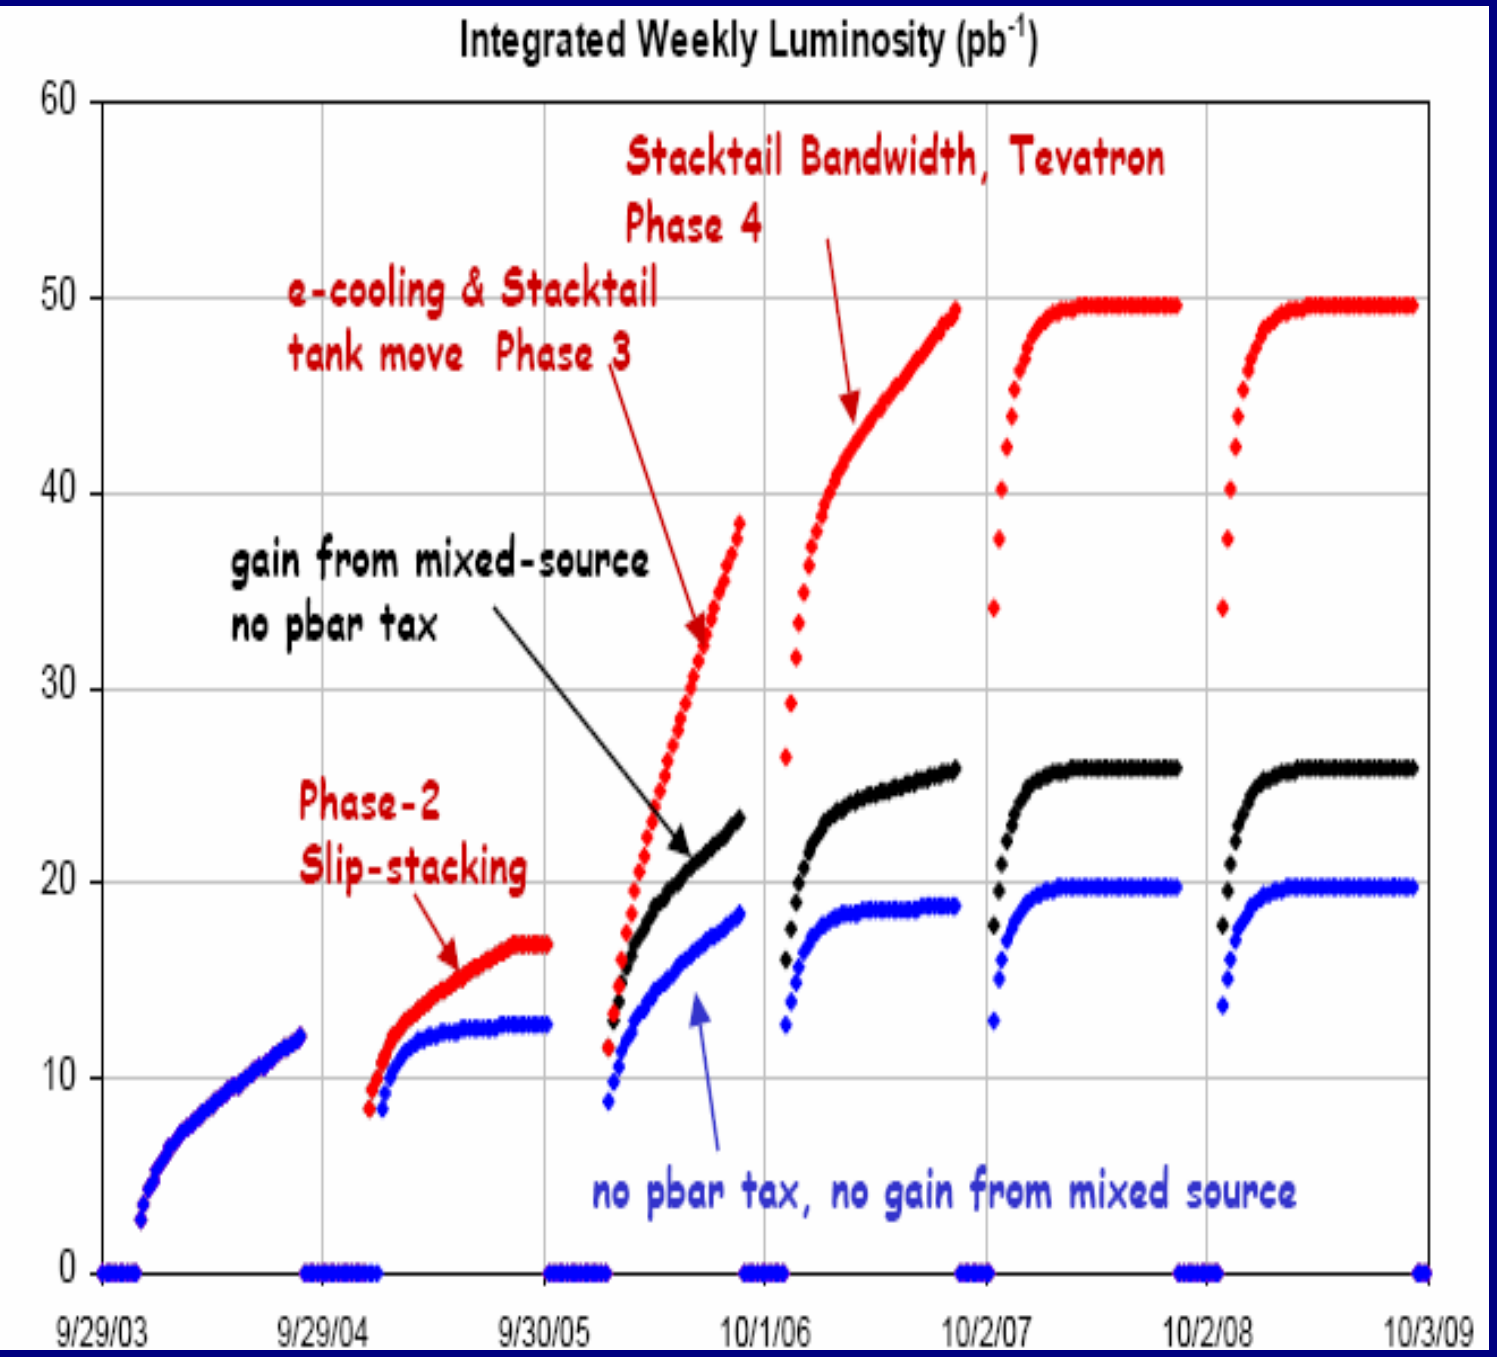

+ FY 2010 ?

Still large factors to be gained over the next few years

Triggers, Cross Sections and Rates

(updated using latest data)

L2 Trigger path	L2 σ@300E30 (nb)	L2 rate@300E30 (Hz)	
CEM18*	260	80	SM
CMUP18*	200	70	
CMX18_JET10*	900	300	
MET15_PEM20	50	15	
WNOTRACK	100	30	
ZNOTRACK	17	5	
TAU10_MET20	50	15	
MET27_CJET25_JET**	450	150	
MET35_JET5	79	24	
4JET_SUMET175	57	17	
BJET15_JET10	140	42	MSSM
CMUP8_TAU*	106	32	
CMX8_TAU*	122	37	
CEM8_TAU*	158	47	
total	**~1100**	**~850**	

*With L1 XFT stereo confirmation (but no L2 yet)

**new trigger as of last table (O. Gonzalez)

=> ~85% of CDF bandwidth

CDF Resources available

	CY 2007	2008	2009
US FTE	222	162	127
Non US FTE	170	135	109
Total US + NonUS	**392**	**297**	**236**
Post Doc's	101	73	53
Students	147	102	77

Collaboration members available in units of FTE

~25% more FTE in CY07 than estimated in 2005

It takes ~100 FTE to Run CDF

CDF Datasets 5/07

period #

run range

dates

#events/lumi

integrated lumi

Data Processing Status

Period	Run-Range	online time	Validation	Release Patch	Dataset	Event numbers / lumi.	Integrated Lumi.	Status	Known Problems
12	237845-241664	01 Apr 07 - 13 May 07	checklist	6.1.1c_f	0j	256M events /185pb-1	2419 pb-1	coming	-
11	233133-237795	31 Jan 07- 30 Mar 07	checklist	6.1.1c_f	0j	369M events /264pb-1	2234 pb-1	half way done	-
10	228664-233111	24 Nov 06- 31 Ja 07	checklist	6.1.1c_f	0i	390M events /280pb-1	1970 pb-1	complete	CEM scale calibration (cache)
9	222529-228596	01 Sep 06 - 22 Nov	checklist	6.1.1c_f	0i	250M events /180pb-1	1690 pb-1	complete	WHA calibration (cache)
8	217990-222426	09 Jun 06 - 01 Sep 06	checklist	6.1.1c_e	0i	335M events / 210pb-1	1510 pb-1	complete	Hot WHA Tower (cache)
7	210012-212133	14 Jan 06 - 22 Feb 06	checklist	6.1.1	0i	~ 105M events / 50pb-1	1300 pb-1	complete	-
6	206990-210011	10 Nov 05 - 14	checklist	6.1.1e	0i	190M events/ 110pb-1	1250 pb-1	complete	-

Luminosity

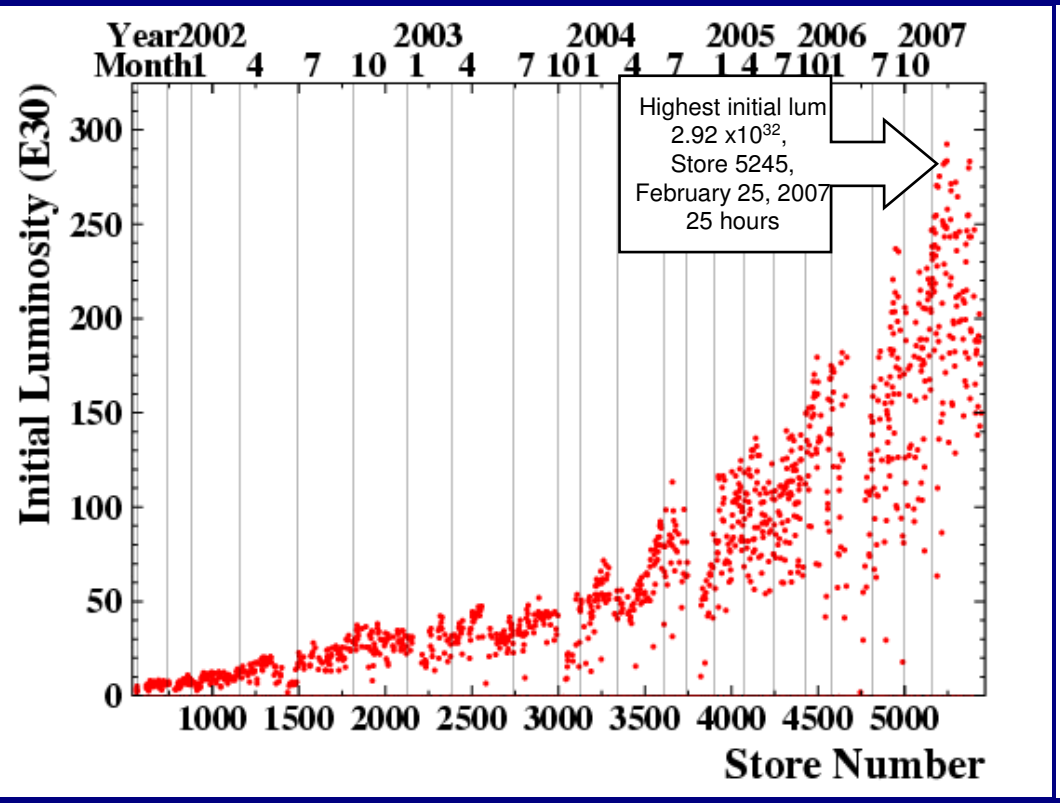
Year2002 2003 2004 2005 2006 2007
Month1 4 7 10 1 4 7 10 1 4 7 1 4 7 10 1 4 7 10
Initial Luminosity (E30)
300
250
200
150
100
50
0
1000 1500 2000 2500 3000 3500 4000 4500 5000
Store Number
Highest initial lum
2.92 x10^32,
Store 5245,
February 25, 2007
25 hours

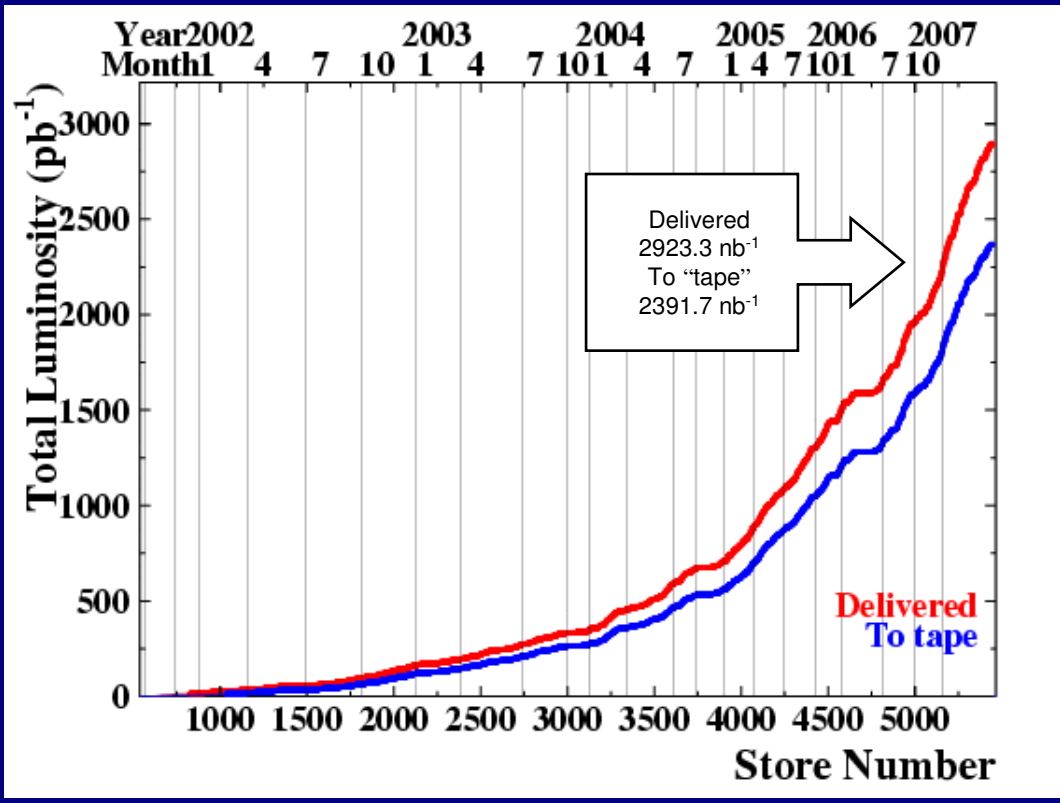
Year2002 2003 2004 2005 2006 2007
Month1 4 7 10 1 4 7 10 1 4 7 1 4 7 10 1 4 7 10
Total Luminosity (pb^-1)
3000
2500
2000
1500
1000
500
0
1000 1500 2000 2500 3000 3500 4000 4500 5000
Store Number
Delivered
2923.3 nb^-1
To "tape"
2391.7 nb^-1
Delivered
To tape

Luminosity Projections with Updated Model Scenarios

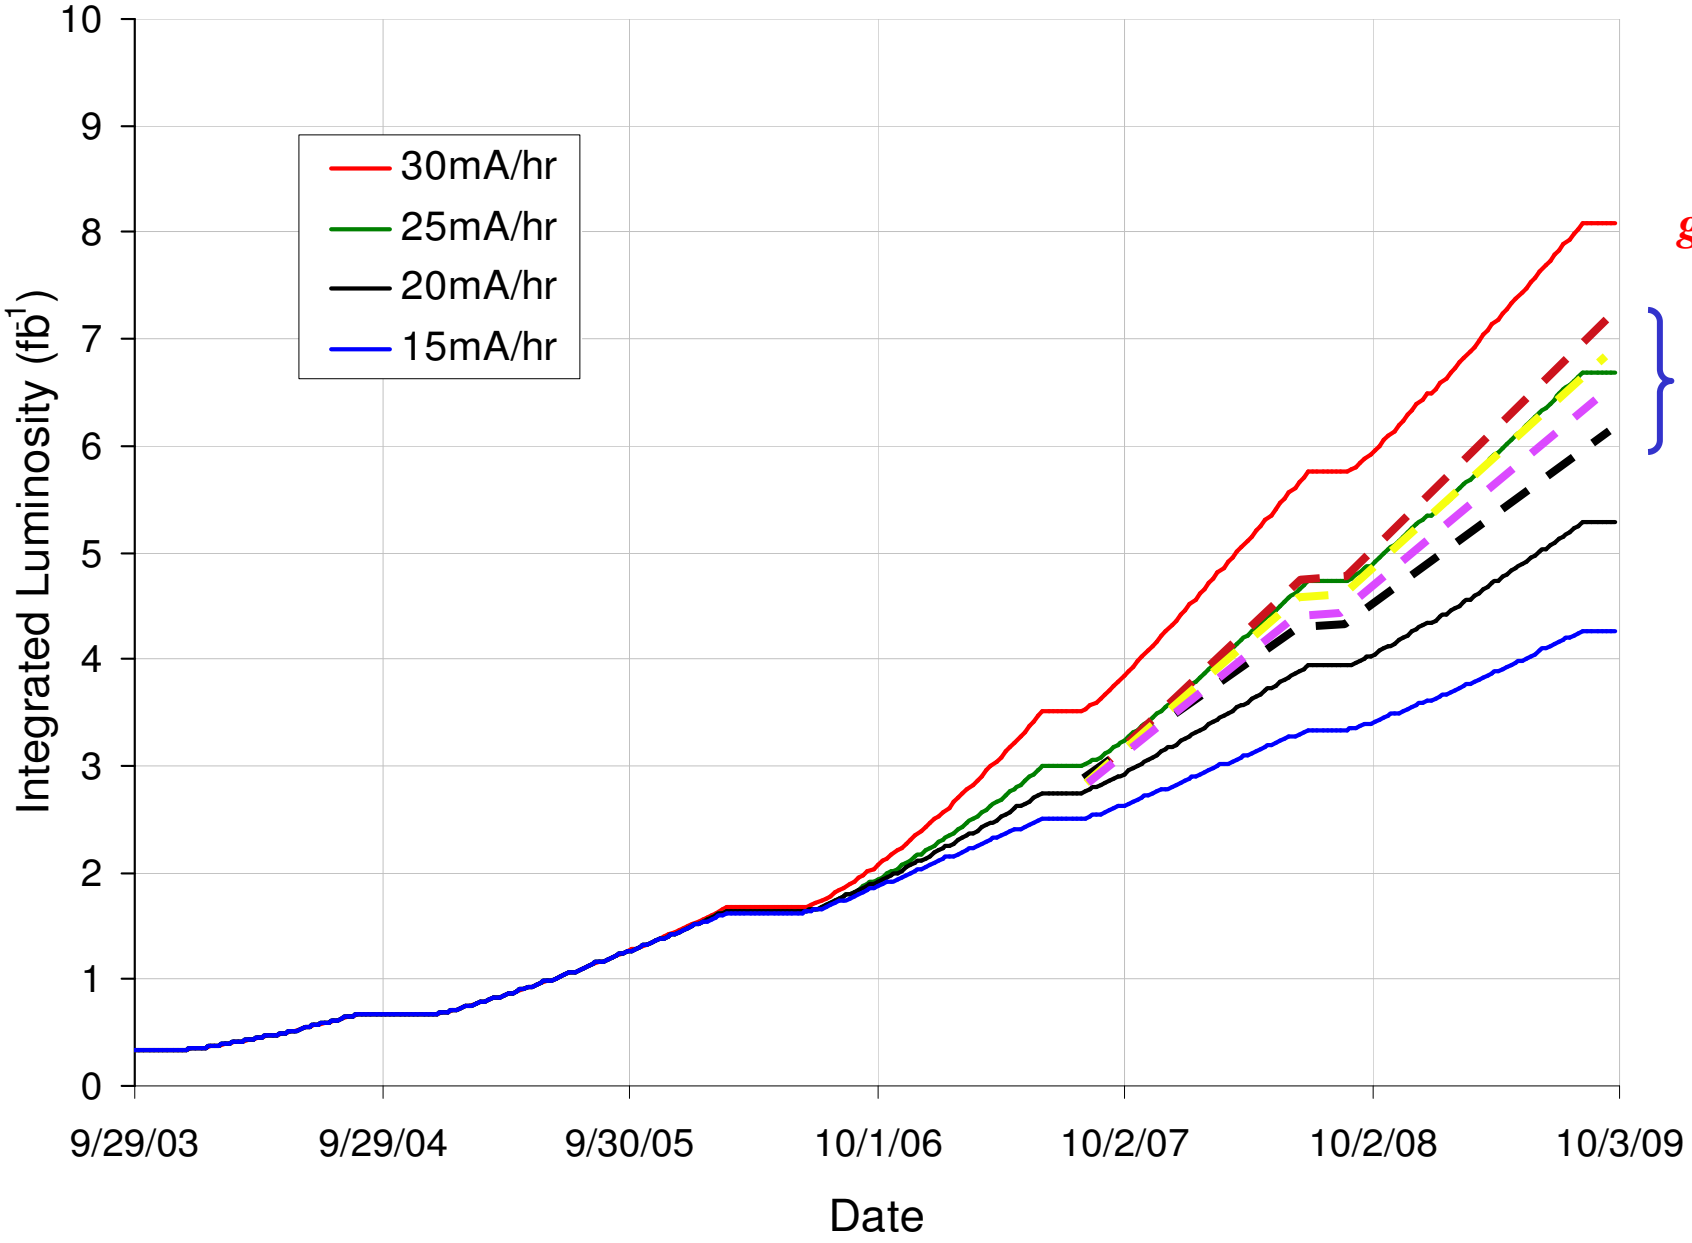

goal

updated model scenarios 1 to 4

≈1 fb^{-1} spread in total

Integrated Luminosity

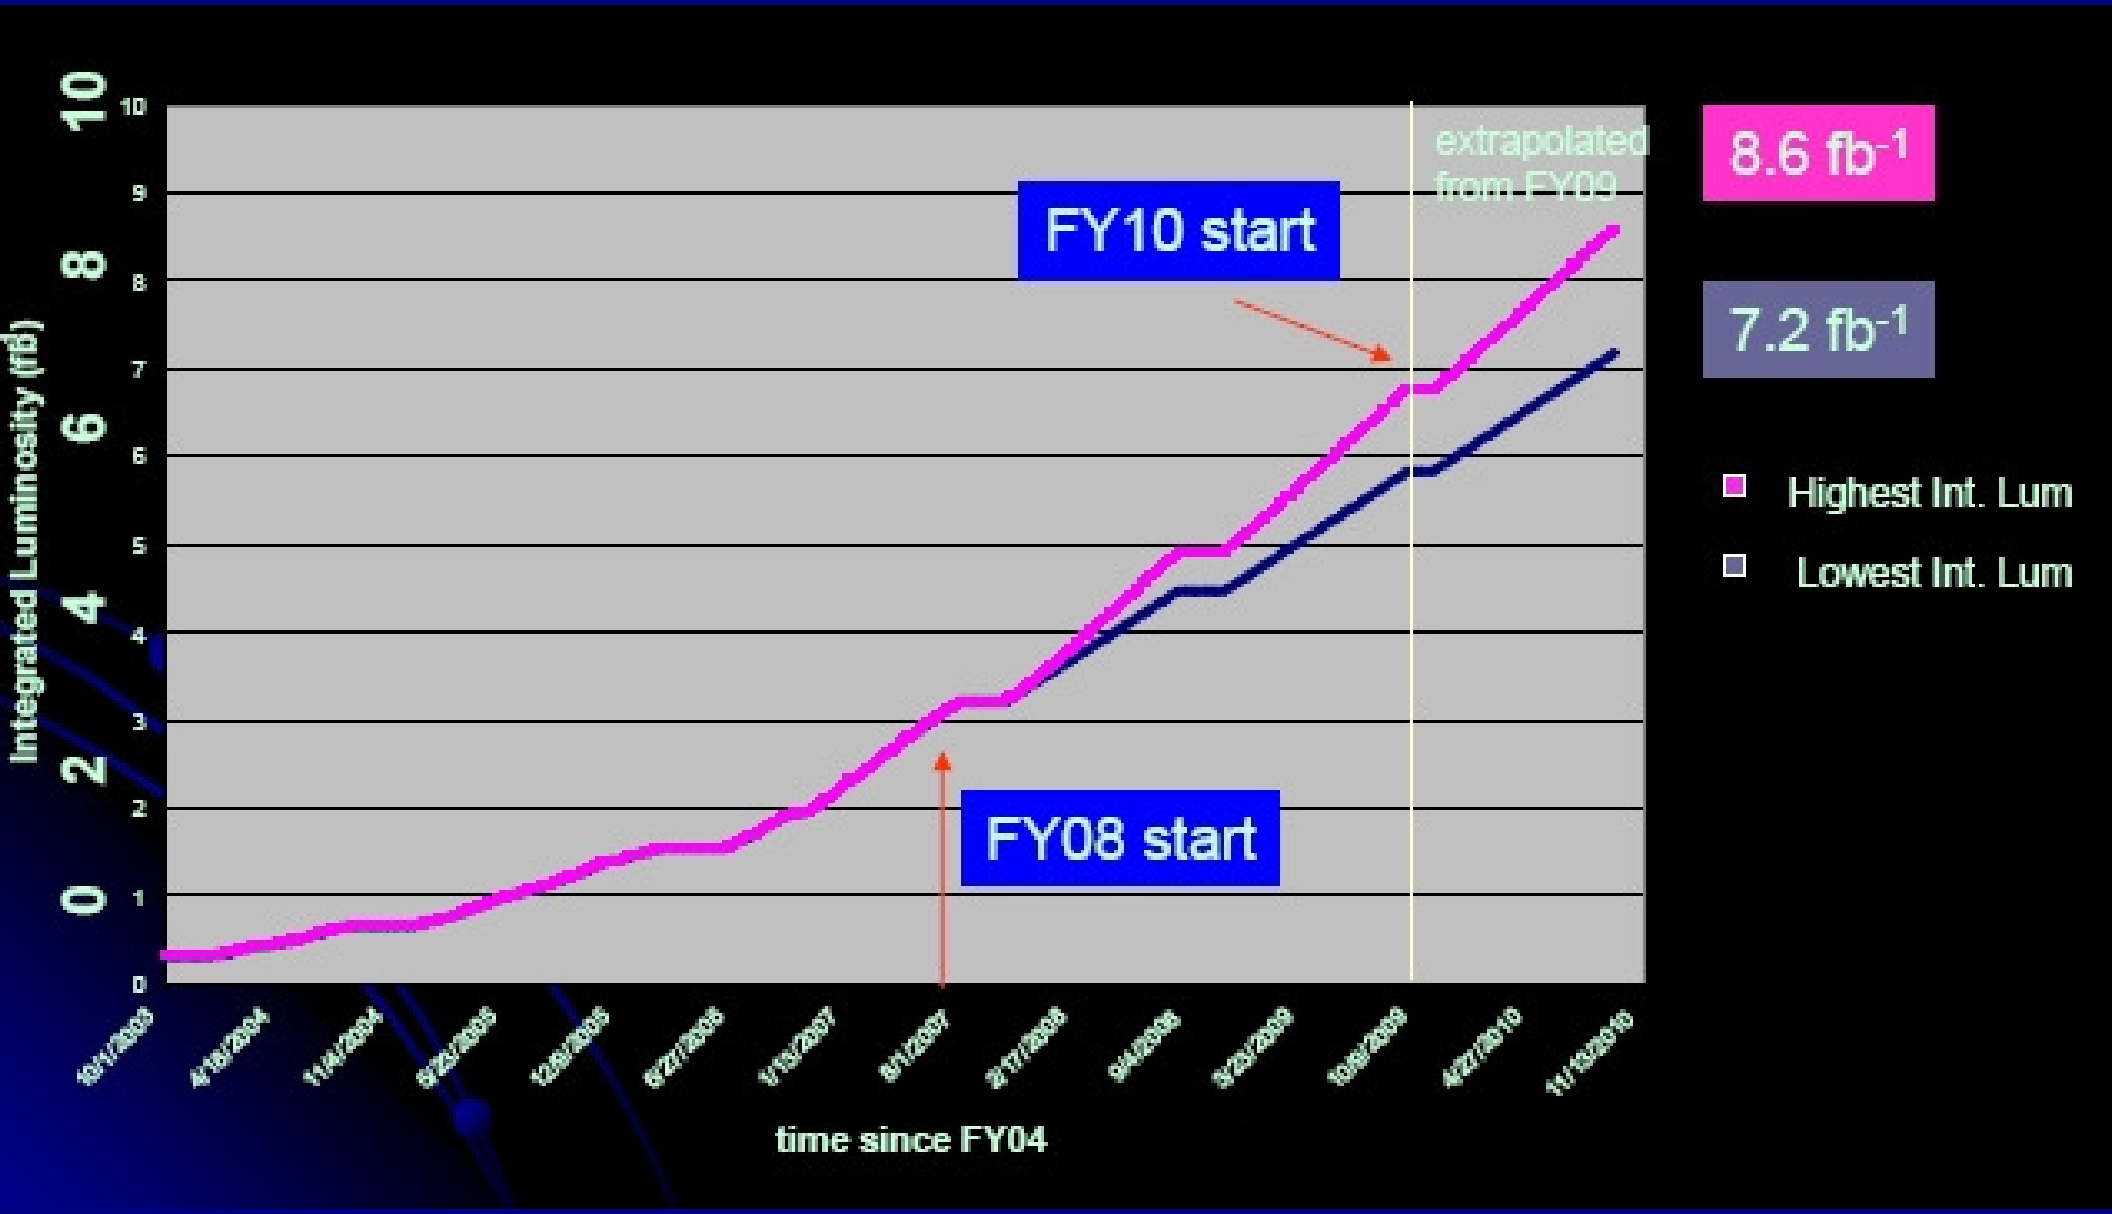

Operations

- **Production processing** (R. Culbertson, E. Gerchtein, R. Harr, B. Jayatilaka, T. Miao, M. Vogel, A. Warburton + calibrators, ntuplers, MC producers)
  - Raw data / ntuple production proceeding on schedule
    - P13 raw data production completed, ntuples almost done
    - Infrastructure / error handling improvements over past few months
      - Processed P13 at record rates (>40 M events/day)
      - Concatenation throughput higher than in the past.
      - Working to further reduce the time for recoveries, clean-up
    - Start P14 after calibration sign-off in about 2 weeks
  - MC production
    - Problem with latest tarball (patch J) delaying P13 MC
      - Expect a resolution within days

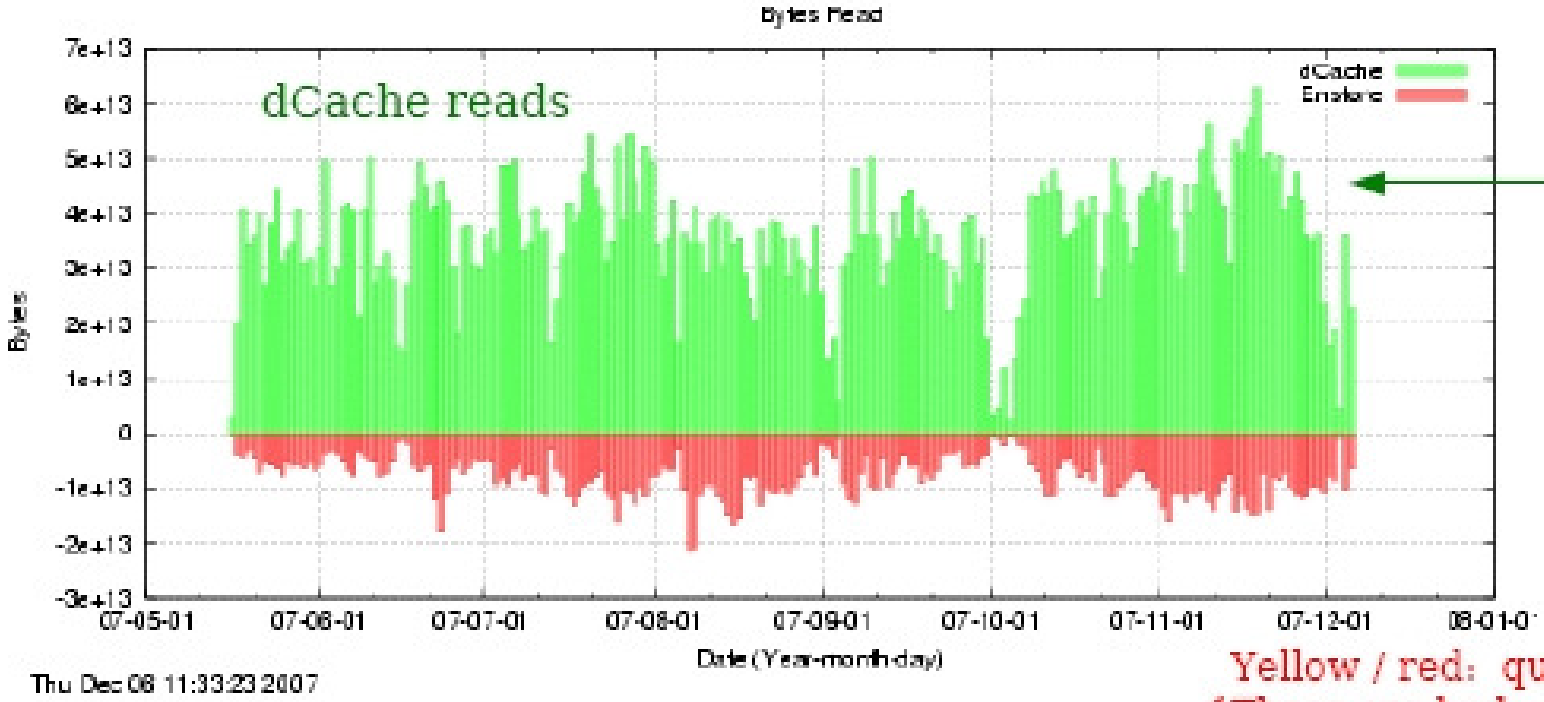

Increased read rate over past month

Yellow / red: queued read / tape restores
These are bad

Manual intervention often required
Re-configurations may fix some cases

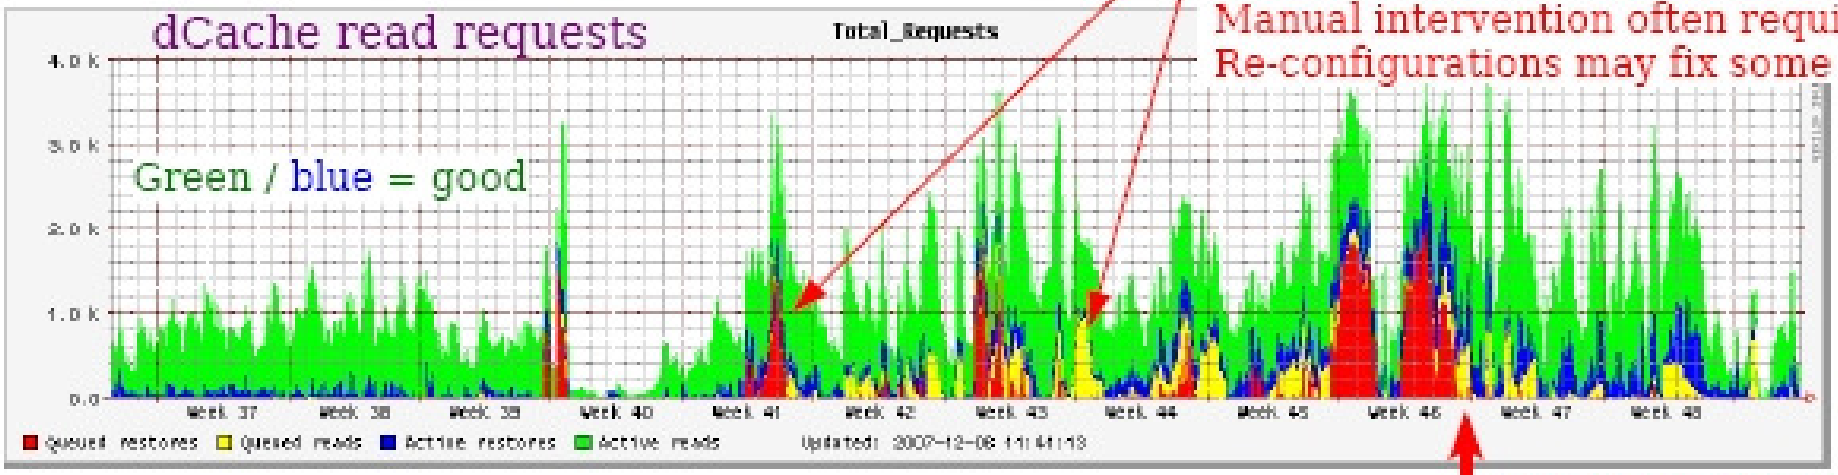

Fixed tape queuing(?)

Computing Time vs Luminosity

- Normalize luminosity curve to unit area in $[50\times10^{30}, 300\times10^{30}]$ cm^{-2} s^{-1}
- Parameterize CPU time curves with 2nd order polynomials
- Convolute the two curves: average CPU time per event

$$\langle t \rangle = \sum_{\text{lumi bins } i} t_i \cdot N_i \cdot \Delta\mathscr{L}_i$$

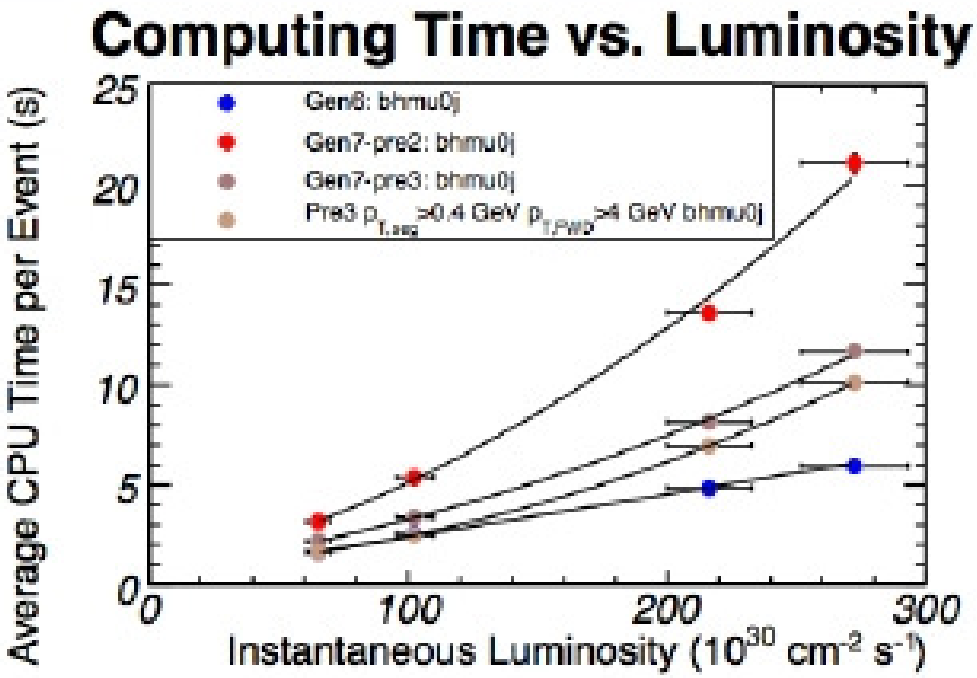

- Results:

Release	Average Time (s)	Ratio to Gen6
Gen6	2.72	1.00
Gen7-pre2	6.69	2.46
Gen7-pre3	4.12	1.52
Gen7-pre3, FWD p_T > 4 GeV/ c	3.42	1.26
Gen7-pre3, Segment p_T > 0.5 GeV/ c	3.33	1.23
Gen7-pre3, FWD p_T > 2, Seg p_T > 0.43	3.26	1.20

Installed Enstore Systems

Enstore provides distributed access to and management of data stored on tape. It provides a generic interface so experimenters can efficiently use mass storage systems as easily as if they were native file systems.

STKEN Enstore System	Mass Storage Production Service for General Fermilab Users
CDFEN Enstore System	Mass Storage Production Service for CDF Run II
DOEN Enstore System	Mass Storage Production Service for D0 Run II
GCCEN Enstore System	Mass Storage Internal Testing/Debugging
Production System's Overall Status	Status for all Production Enstore systems

Total User Data on Tape (Cdfen, D0en, Stken) :	7996.090 TB

http://rexganglia1.fnal.gov/

Available ganglia pages:

Farms (CDF, D0, and GP)

CDF Offline

CDF Online (requires login)

D0 Offline

D0 Online

MINOS

Cluster Report for Wed, 6 Feb 2008 05:43:46 -0600

Get Fresh Data

Metric: load_one Last: hour Sorted: descending

Physical View

Farms Grid > CDF Farm > --Choose a Node

Overview of CDF Farm

CPUs Total:	**150**
Hosts up:	**73**
Hosts down:	**3**

Avg Load (15, 5, 1m): **46%, 49%, 48%**

Localtime: **2008-02-06 05:43**

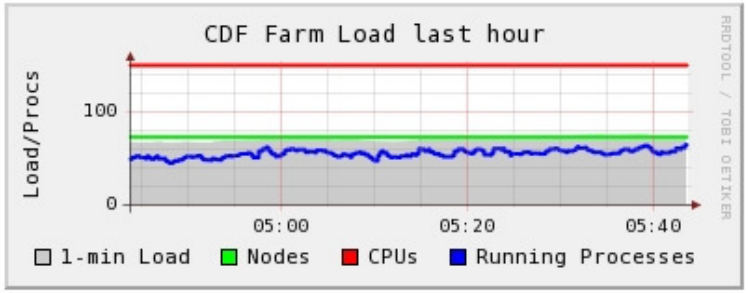

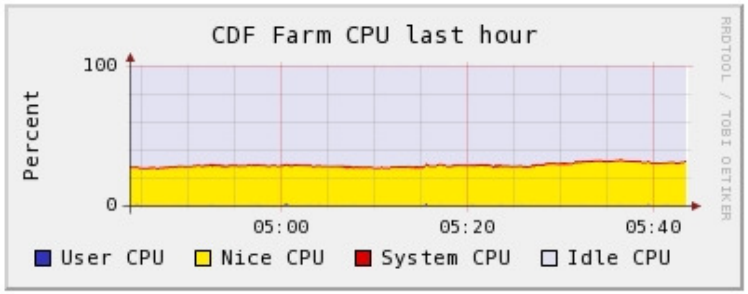

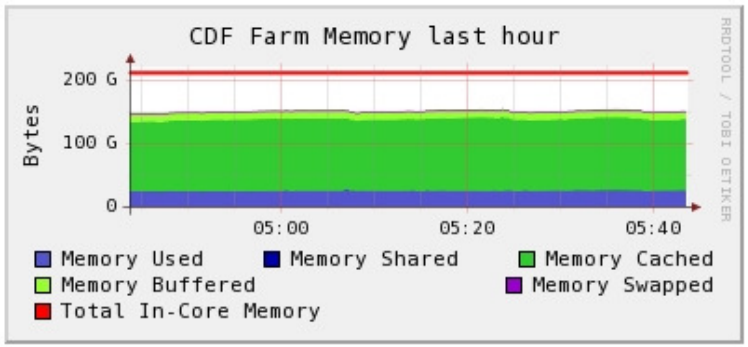

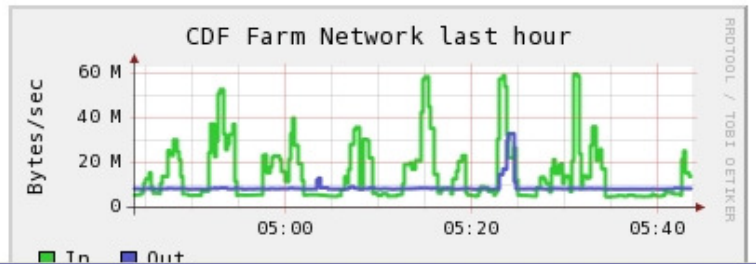

Pie Chart

Fermi National Accelerator Laboratory

FEF Faultlog

D0	CDF	FEF	GP Grid	MiniBoone
MINOS	MIPP	SCIBOONE		

Overview

- Ganglia thinks 16 nodes are down. (3275 up) View all.
- Of those 16 machines, 0 outages have been acknowledged*. View all.
- 908 nodes are not reporting to Ganglia.
- 4199 total nodes found in SYSADMIN database.

Outages

- 783 entries in faultlog with a recorded outage duration.
- Of those, the shortest outage was for the host FND0749, with a duration of 3 minutes.
- The longest outage was for D0OL95, and lasted for 153 days.
- Average outage is 9 days.

Queries

- See activity for past seven days.
- Generate a list of machines not reporting to syslogDB (slow link, takes ~2 minutes to load)
- Locate a single node: (partial names accepted)

Search

Update Frequency

- Ganglia data updates once per minute.
- Hardware calls update once per hour.
- Cluster information updates once per day, early in the morning.

* Acknowledged is defined as a faultlog entry that is more recent than the last time the machine reported to ganglia.

Towards GRID

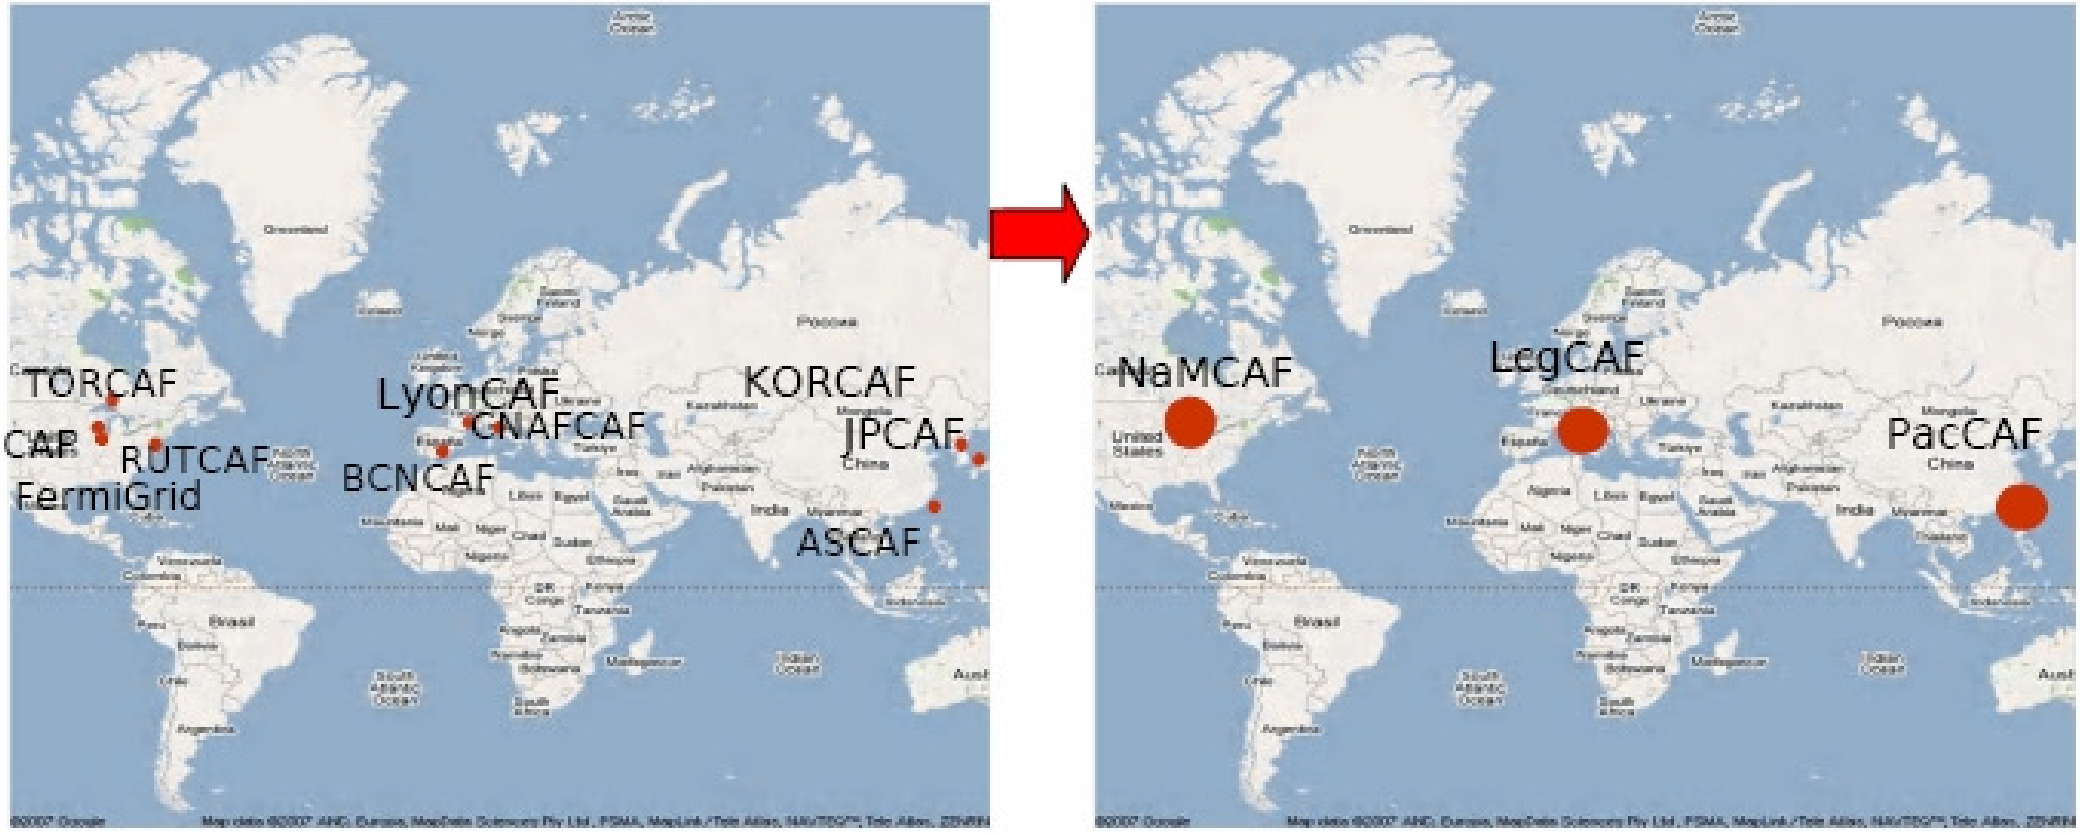

Eliminate all but three grid submission portals: NAMCAF, LcgCAF, PacCAF.

Migrate all existing systems accordingly. (May keep FermigridCAF+CNAF for data access.)

Current CDF Dedicated Resources

Current Resources [*]			
Cluster Name and Home Page	Monitoring and Direct Information Links	CPU (GHz)	Disk space (TBytes)
Original FNAL CAF	queues, user history, analyze, ganglia, sam station, consumption	1000	370
FNAL CondorCAF (Fermilab)	queues, user history, analyze, ganglia, sam station, consumption	2200	(shared w/CAF)
CNAFCAF (Bologna, Italy)	queues, user history, analyze, resources, network, sam station, datasets, consumption	480	32
KORCAF (KNU, Korea)	queues, user history, ganglia, sam station, datasets, consumption	178	5.1
ASCAF (Academia Sinica, Taiwan)	queues, user history, ganglia, sam station, datasets, consumption	134	3.0
SDSC CondorCAF (San Diego)	queues, user history, analyze, ganglia, sam station, datasets, consumption	380	4.0
HEXCAF (Rutgers)	queues, cpu, sam station, datasets, consumption	100	4.0
TORCAF (Toronto CDF)	queues, user history, analyze, ganglia, disk status, sam station, datasets, consumption	576	10
JPCAF (Tsukuba, Japan)	queues, user history, ganglia, sam station, datasets, consumption	152	10
CANCAF (Cantabria, Spain)	queues, user history, ganglia, sam station	50	1.5
MIT (Boston, USA) (MC only)	queues, user history, analyze	322	3.2
	Current Totals (*):	5572	448

http://www-cdf.fnal.gov/internal/fastnavigator/fastnavigator.html (2006/Aug)

Usage of Dedicated Farms: all dCAFs

rtime from September 06 up to now

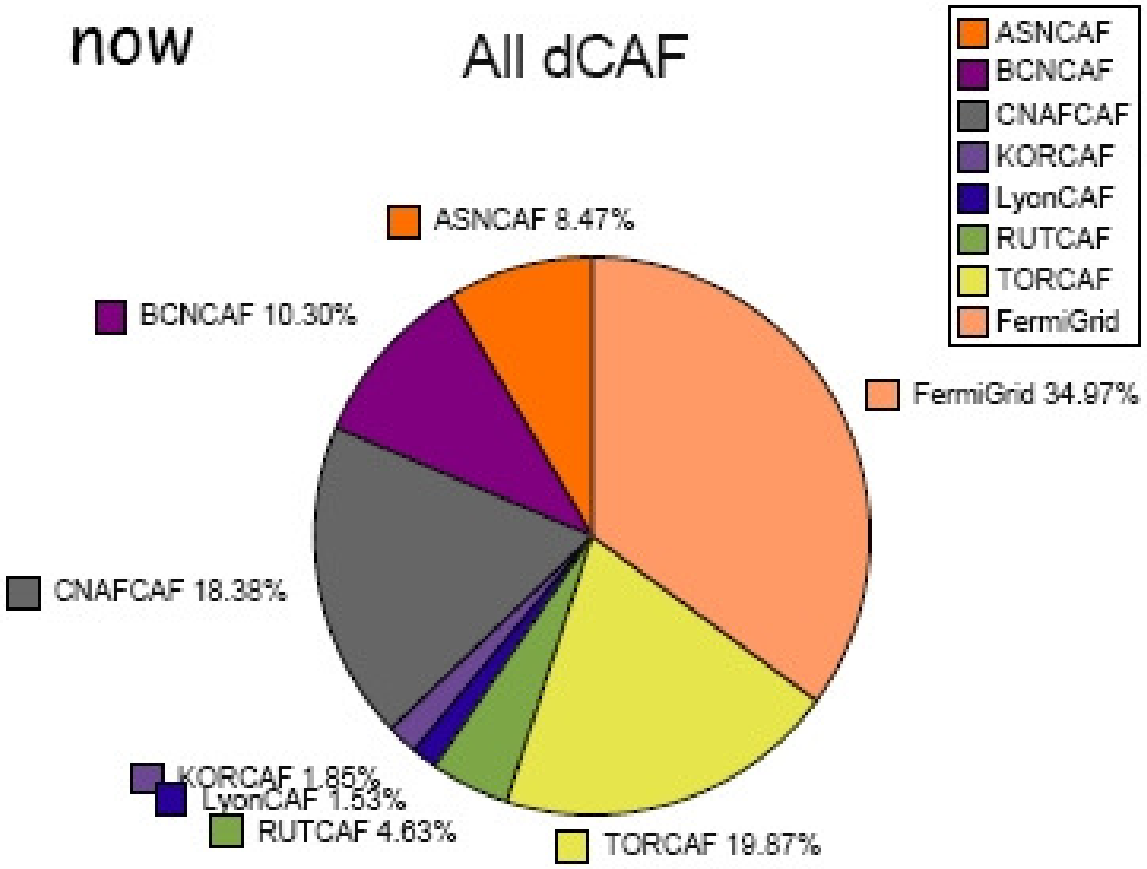

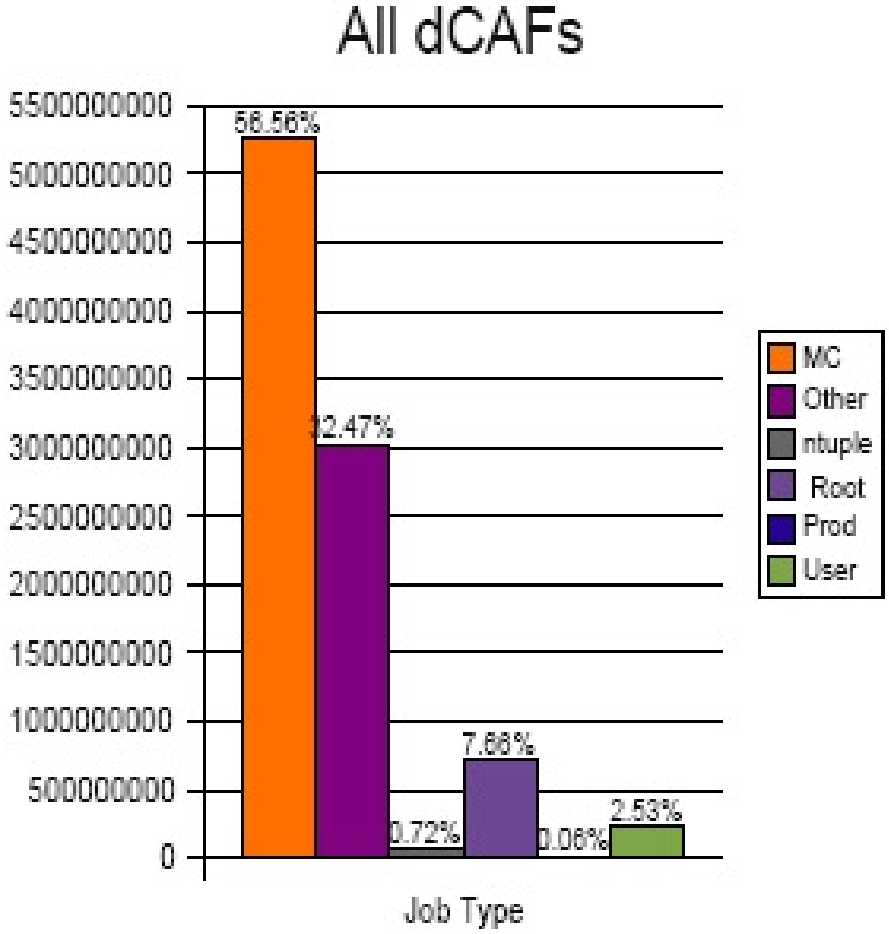

Usage of Dedicated Farms: CAF

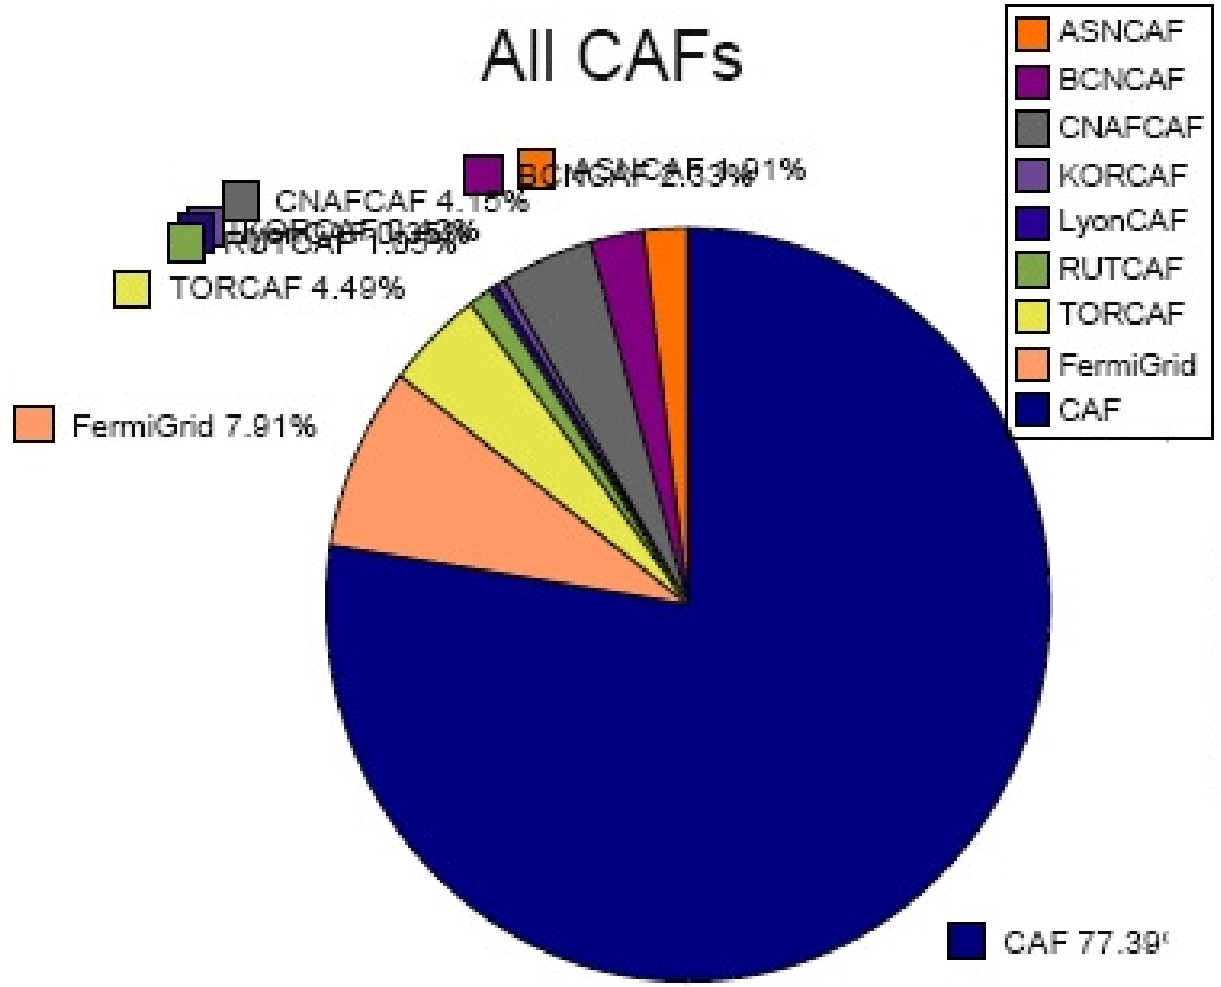

CAF is a great system but too overloaded!

Policy: "Use CAF only for data"

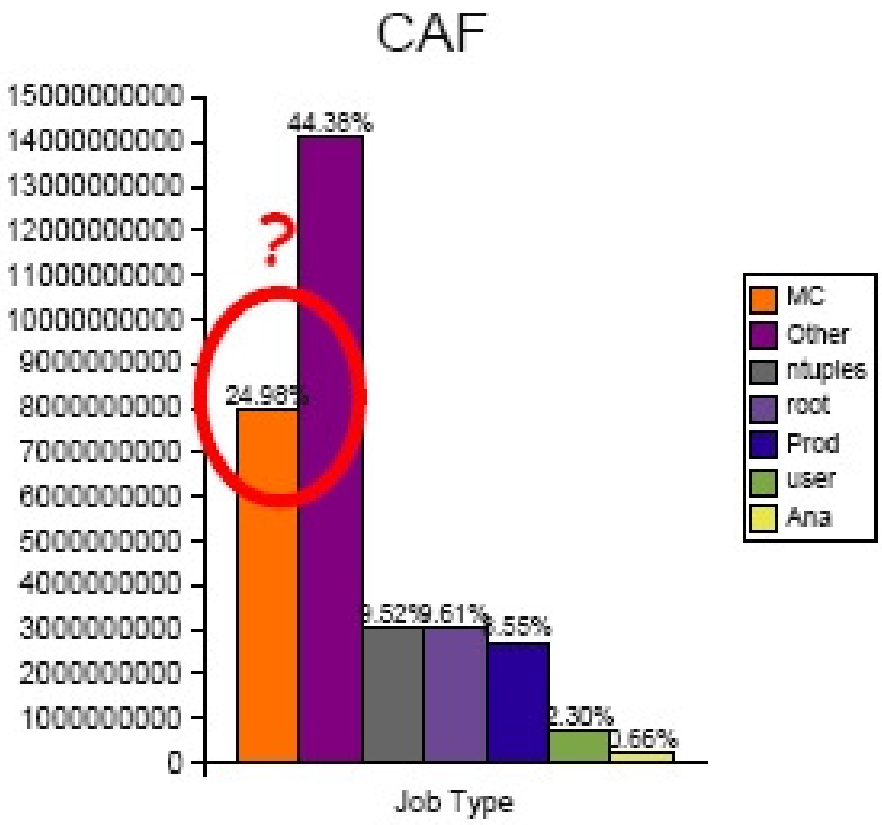

Being overloaded contributed to the recent failures

Useful Info

Proposals / Plans & PO 2004 – 2007

Age and warranty status can mostly be determined from
http://appora.fnal.gov/equipdb/equipDetails.html

general rule,
any system greater than 3 years old is no longer under warranty

Future Challenges

- Higher instantaneous luminosity
  - Larger events, slower reconstruction, tracking more difficult, need more CPU per event
- Higher integrated luminosity and higher data taking rate
  - Larger data samples
    - Need more processing power
    - Need more storage
- Migration of physicists to LHC experiments
  - Human resources for operations are shrinking
- FY2010 Running has been proposed