Computing model and budget update

Rick Snider
Fermilab

Outline

- Computing model
- Requirements and budget
- Future budget planning issues

Spring CDF IFC Meeting
April 29, 2006

CDF computing model

- Main issues relevant to computing plan and budget
  - Increases in raw data logging rate
    - Final planned increase to 60 MB/sec in FY2006 (Q4)
  - Limited Fermilab computing budget
    - Cannot meet demand locally
    - Need to reduce costs, increase operational efficiency
    - Continued slow pace of CPU speed increases
  - Evolving grid infrastructures, access policies
    - Changing how we deploy, exploit resources
  - Basic data flow, computing systems, requirements the same
    - Many details are different

CDF computing model

- Major hardware systems
  - Tape archive + disk cache
  - Production farm (uses CAF infrastructure)
    - Performs full reconstruction of all data
  - CAF
    - Ntuple production, user analysis, MC generation
    - 2 dedicated farms + 1 CDF-purchased farm in Fermigrid (OSG) + opportunistic use of other Fermigrid computing elements
  - dCAFs
    - 10 remote sites — 4 in shared pools accessed via glide-ins
  - Databases + networks + static “project” disk

CDF computing model

- Major hardware sub-systems
  - Tape archive + disk cache
  - Production farm (uses CAF infrastructure)
    - Performs full reconstruction of all data
  - CAF
    - Ntuple production, user analysis, MC generation
    - 2 dedicated farms + 1 CDF-purchased farm in Fermigrid (OSG) + opportunistic use of other Fermigrid computing elements
  - dCAFs
    - 10 remote sites — 4 in shared pools accessed via glide-ins
  - Databases + networks + static “project” disk

New in past 6 months

CDF computing model

- Data flow
  - Raw data written to tape
  - Measure final calibrations (4 weeks of data per cycle)
    - Process about 30% of data stream on production farm and CAF
  - Full (final) reconstruction on production farm
    - Input read, output written to tape
  - Centralized ntuple production
    - Two of three ntuples produced on production farm
  - User analysis of ntuples and production output on CAF/dCAF
  - Monte Carlo generation/reconstruction on remote dCAFs
  - Final analysis on user desktops / institutional machines

CDF computing model

New in past 3 – 12 months

- Data flow
  - Raw data written to tape
  - Measure final calibrations (4 weeks of data per cycle)
    - Process about 30% of data stream on production farm and CAF
  - Full (final) reconstruction on production farm
    - Input read, output written to tape
  - Centralized ntuple production
    - Two of three ntuples produced on production farm
  - User analysis of ntuples and production output on CAF/dCAF
  - Monte Carlo generation/reconstruction on remote dCAFs
  - Final analysis on user desktops / institutional machines

Better consolidation

CDF computing model

- Data flow
  - Raw data written to tape
  - Measure final calibrations (4 weeks of data per cycle)
    - Process about 30% of data stream on production farm and CAF
  - Full (final) reconstruction on production farm
    - Input read, output written to tape
  - Centralized ntuple production
    - Two of three ntuples produced on production farm
  - User analysis of ntuples and production output on CAF/dCAF
  - Monte Carlo generation/reconstruction on remote dCAFs
  - Final analysis on user desktops / institutional machines

The main points

Computing requirements

- Computing requirements model (presented last fall)
  - Resource demand depends upon:
    - Dataset volume
    - Data taking rate
  - Dataset volume dependence dominates
    - Adopted “proportional” model
  - Proportional model: scale FY2005 resources
    - FY200x demand = (FY2005 resources) $\times$ (FY200x/FY2005 data vol)
  - Estimated using “proportional” model

Computing model input parameters

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

No changes since Fall IFC

* 40 MB/s until Q3

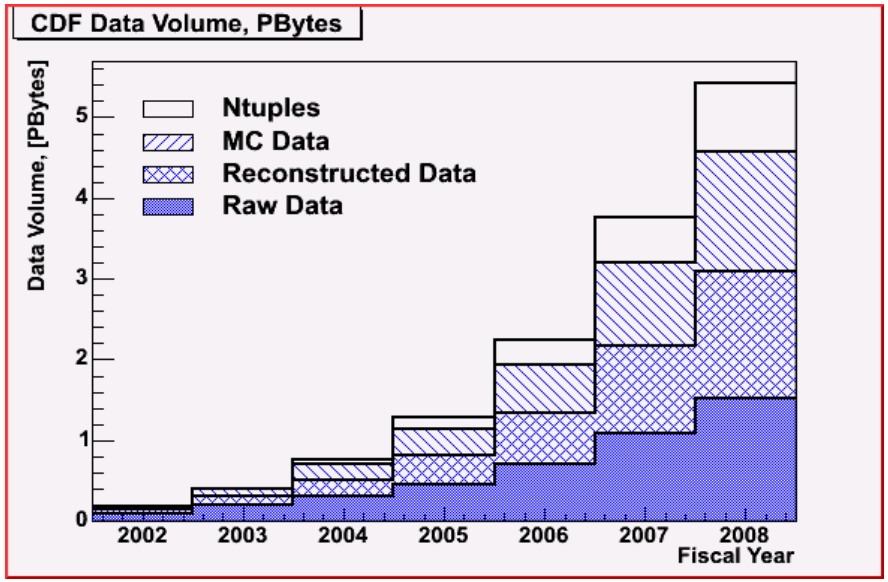

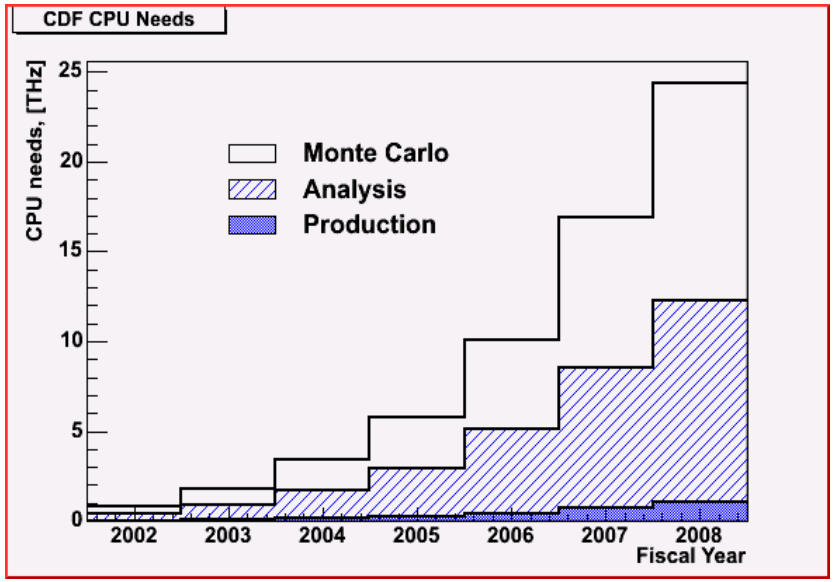

Projected computing requirements

Fiscal Year	2006	2007	2008
CPU needs (THz)*	10	18	24
Disk volume (PB)†	0.7	1.0	1.4
Data volume on tape (PB)	2.2	3.8	5.7

No changes since Fall IFC

* Shared between CAF and dCAF's
† Hosted at Fermilab

Requests to IFC

- Standing request to IFC for past 3 years
  - Provide 50% of CPU requirement
    - Includes contributions located on-site
    - Total CPU requirement scales at about 60% per year
  - No disk beyond that on worker nodes (except for analysis centers)
    - May eventually need ~1 TB for local buffering
- Status of physics analysis center proposal from fall
  - Discussing prototype analysis center at CNAF
  - Will decide how to proceed after some experience

Current inventory

		Actual		Requirements		
	Fiscal Year	2005	2006	2006	2007	2008
CPU (Thz)	Fermilab	3.4	4.8	5.0	8.8	12
	On-site contributions†	1.5	1.7	1.7	1.7	1.7
	Remote†	2.3	2.3	3.3	7.1	10
	Total	7.2	8.8	10	18	24
Disk (PB)	Fermilab	0.2	0.6	0.5	0.9	1.3
	On-site contributions	0.2	0.2*	0.2	0.1	0.1
	Remote	0.1	0.1?			
	Total	0.5	0.9?	0.7	1.0	1.4
Volume on tape, April 2006 (PB)		1.3	1.5	2.2	3.8	5.7

* 50% retires at end of FY2006

† Compare sum to Fermilab

Setting a lower bound

- Two extreme cases in modeling computing requirements
  - Proportional model:
    - Requirements scale with total data volume
  - Incremental model:
    - Requirements scale with data logging rate
- Real life is somewhere between these (more on this later)
- Use incremental calculation to set approximate lower bound on computing needs
  - Assume FY2005 – FY2006 resources were approximately correct size

Requirements in a purely incremental model

(an approximate calculation)

		Actual		Requirements			
	Fiscal Year	2005	2006	2006	2007	2008	
CPU (Thz)	Fermilab	3.4	4.8	5.0	6.0	6.0	
	On-site contributions†	1.5	1.7	1.7	1.7	1.7	
	Remote†	2.3	2.3	3.3	4.3	4.3	
	Total	7.2	8.8	10	12	12	±1 THz
Disk (PB)	Fermilab	0.2	0.6	0.5	0.7	0.7	
	On-site contributions	0.2	0.2*	0.2	0.1	0.1	
	Remote	0.1	0.1?				
	Total	0.5	0.9?	0.7	0.8	0.8	
Volume on tape, April 2006 (PB)		1.3	1.5	2.2	3.8	5.7	

* 50% retires at end of FY2006

† Compare sum to Fermilab

Projected cost estimates

- Purchase FY200x hardware with FY(200x – 1) funds

Projected computing requirements

Fiscal Year	2006	2007	2008
CPU needs (THz)*	10	18	24
Disk volume (PB)†	0.7	1.0	1.4

* Shared between CAF and dCAF's
† Hosted at Fermilab

Projected total computing cost ($M)

Fiscal Year	2006	2007	2008
Total cost	2.9	2.8	2.5

Current Fermilab computing budget

- Current Fermilab budget

FY2006	Equipment ($M)	Operating* ($M)
Request	1.47	1.0
Budget	1.27	1.01
Spent	–	0.24

* Tapes + maintenance + …

  - Additional $122k contributed from Japan in FY2006
- Working on FY2006 procurements
  - Will basically follow plans presented in Oct.

Future budget planning

- Important considerations in planning for the future
  - Evolution of the computing model
  - Expected changes in the available human resources

Changes to requirements model

- Evolution of the computing model
  - Moving toward much more incremental computing model
    - “One-pass” production
      - At most one large-scale re-processing foreseen
    - Centralized ntuple production fully incremental
    - Run-dependent MC production
    - Moving toward single dataset: delete all but current from tape

    Reduces operational load and improves resource utilization
  - Grid developments (Igor's talk)
  - Working on other strategies to improve resource utilization

Changes to requirements model

- Need to include some effects in budget projections, e.g., incremental computing
  - Have better measures of what we do with our computing

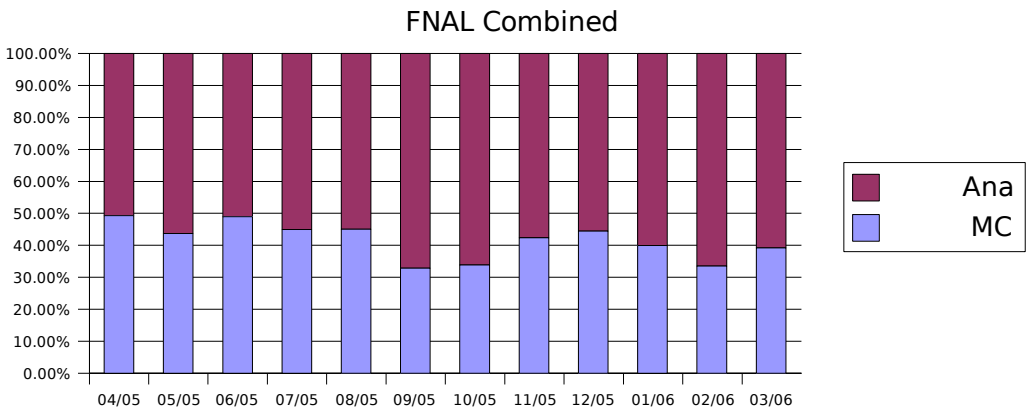

  - Introduce incremental components explicitly rather than assume they are proportional
    - Reduces expected demand relative to non-incremental model

Budgeting with declining effort

- Changes in available FTE profile
  - Run II Task Force report
    - FTE's on CDF decline 25% – 30% by 2007, about 50% by 2009
  - What is impact upon computing plan?
- Strategy to deal with decline in available effort
  - Consolidate computing, DH technologies with other expts
    - e.g., LCG / OSG compatibility
    - Adapt rather than develop whenever possible
  - Streamline, better automate operational procedures
    - Calibrations, production farm, etc.

This work is already in progress

Budgeting with declining effort

- Address human resource issue in budget estimates
  - Separate demand/budget estimates into two components
    - Requirements for 10 “core analyses” identified by R2TF
      - Expect effort to be available for this in all scenarios
    - Requirements for balance of program
  - Much computing is common to many analyses
    - Event reconstruction
    - Ntuple creation
    - Standard MC datasets
  - Computing requirements may not fall dramatically

Will complete these estimates for fall reviews

Summary

- The success of CDF physics program owes much to the strong support and contribution of the Computing Division and IFC to the CDF computing project.
- As CDF continues to evolve the computing model and adapt to changing circumstances, seek to:
  - Always reduce costs and required effort — do things better
  - Reflect these changes in the way we estimate our computing requirements
  - Spend the resources wisely

  Must have a sustainable system to continue producing physics through the end of the experiment

The end

Backup slides

Equipment budget details

Projected total computing cost ($M)

Fiscal Year	2006	2007	2008
CPU*	2.0	1.9	1.8
Disk	0.4	0.3	0.2
Tape drives + library	0.2	0.4	0.3
Interactive computing	0.1	0.02	0.02
Databases	0.03	0.03	0.03
Miscellaneous	0.1	0.1	0.1
Networking	0.2	0.1	0.1
Total cost	2.9	2.8	2.5

* Cost shared between Fermilab and remote institutions

Fermilab FY2006 budget details

- Equipment budget for FY2006
  - CPU: $500k
  - Disk: $285k
  - Tape drives + library: $200k
  - Networking: $100k
- Operating budget (M&S)
  - Tapes: $300k
    - $90k spent so far
  - Maintenance: $390k
- FY2006 procurement strategy under discussion

CDF computing model

- Major software systems (with substantial ops/dev load on CDF)
  - SAM data handling system
    - dCache disk caching managed by SAM
  - Databases
    - DB application software
    - Frontier
      - Used to distribute calibration DB data
  - CAF/dCAF infrastructure
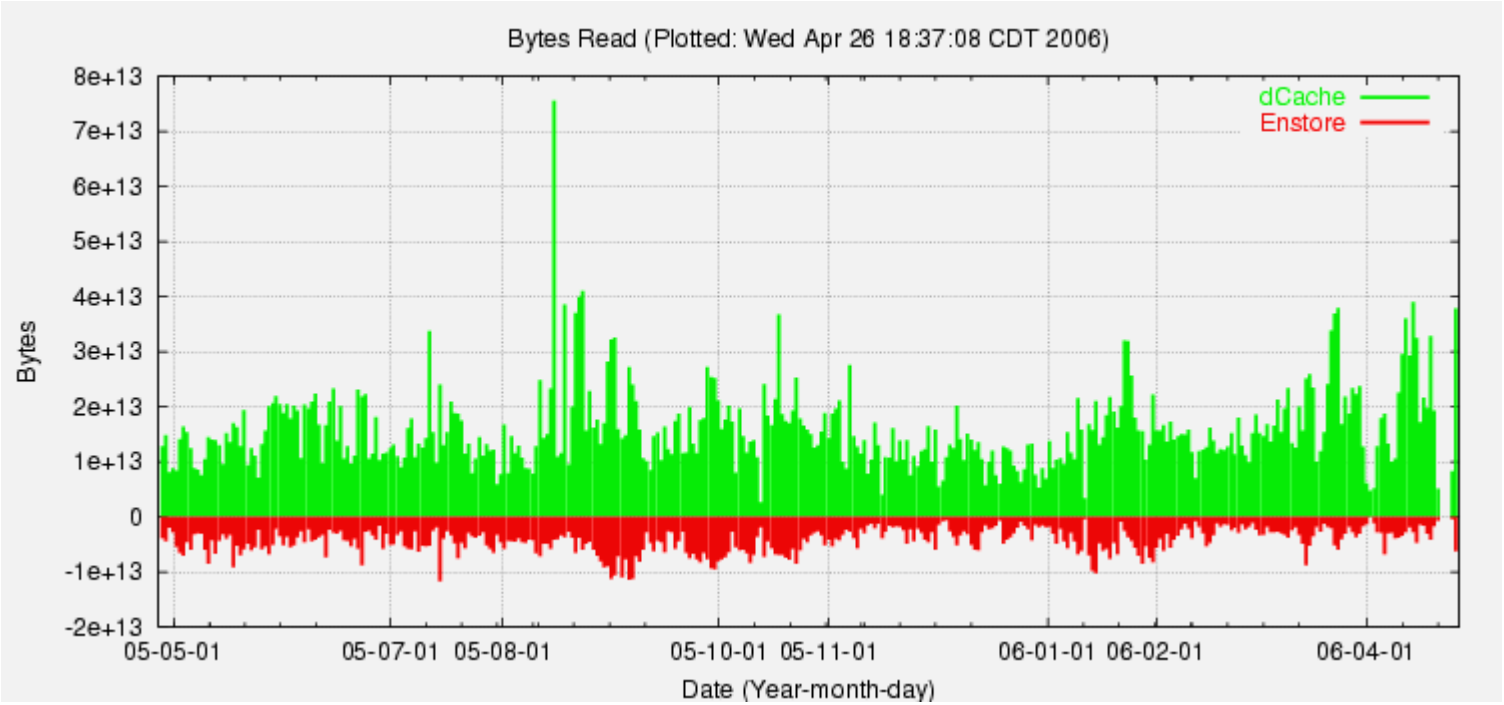
Data handling system throughput