CDF Data Preservation Strategy

Rob Roser
Fermilab

(Adopted and updated version of Rick Snider's talk last year)

Outline

- CDF data processing
- Goals for CDF effort
- Retaining analysis capabilities
- Resource requirements
- Summary of issues

Current Status

- Status at last workshop
  - Management / lab level discussions regarding expectations, needs
    - General agreement that computing capabilities needed to be retained after data-taking ended
    - Initial discussions with collaboration at large
    - No specific targets, requests, or policies defined

- Activity within CDF since last Workshop
  - Several Collaboration Meetings
  - Agreed on some basic definitions and targets
  - Defined some plans for post-data-taking processing
  - Working hard to get current system documented, stable, utilizing latest releases and versions of commercial software etc.

Data processing activities

- Two major classes of processing in CDF computing model
  - Data / MC production
    - Includes
      - Primary event reconstruction from raw data (local CPU resources)
      - Ntupling (initial data reduction)
        - *Includes additional reconstruction steps*
      - Large-scale Monte Carlo simulation (remote CPU resources)
    - During data taking
      - Completely centralized / coordinated processing activity
      - Predictable, easily calculable demand: scales with data logging rate
    - ...After data taking
      - Demand not as easily predicted, driven by physics requirements
      - Currently expect at least one major re-processing of raw data after run ends
        - *May include all or significant fraction of data*
      - Major MC demand will decline relatively quickly
      - Ntupling re-processings tend to occur more frequently than for raw data

Data processing activities

- Two major classes of processing in CDF computing model (cont'd)

  Analysis

  - Includes a variety of activities, including but not limited to:
    - Analysis specific reconstruction, data reduction, MC simulation
    - Ntuple* analysis
    - Pseudo-experiments, multivariate calculations, systematic error scans, statistical calculations
    - Etc., etc.,... The details change with time...
  - De-centralized, largely uncoordinated activity
    - Code and specific computations required can change rapidly
    - Demand varies widely with availability of data + conference schedule + other outwardly random forces
    - Local CPU resources for data-intensive processing. Prefer remote resource for all else

* ntuple = compressed dataset readable/browsable with root

Data preservation and long-term analysis

- Goals of CDF's Data Preservation and long term analysis planning
  - ❖ To allow CDF to continue data analysis for long period after data-taking ends
  - ❖ What does that mean?
    - ❖ Long Term means 5 years minimal
    - ► In this context data analysis means the ability to perform a full analysis including generation of new MC signal samples etc.
    - ► Preserve required data
    - ► Retain required (current) capabilities or develop replacements for at least those 5 years…
    - ► Operate required systems

For How Long?

- At least two time scales of potential interest
  - ❖ Period of “active” analysis after data-taking ends (guessing 5 years)
  - ❖ Period after the collaboration disbands or no longer actively analyzing data – this time period is yet to be determined and in all likelihood will depend on LHC’s success/findings
    - ❖ IF LHC finds a dessert in terms of new physics in the next two years – our data is perhaps far less important than if a wealth of lower energy new physics is found

Currently considering only period of active analysis

  - ❖ Most interest from collaboration, so most likely to receive effort
  - ❖ A prerequisite to answering any longer-term questions is to solve this case
  - ❖ Does not preclude addressing the very long-term
    - ❖ We don’t have $$$ or manpower currently to look longer term – hard enough for us just to do this

- Policy: assume active analysis will last a min. of 5 years
  - ❖ Choice based on combination of experience from previous experiments (including CDF) and on existing analysis plans and expectations

Anticipated changes (challenges) after data-taking

- What can we expect to happen in the 5 years after data taking
  (based roughly on the last 5 years)
  - ❖ All lab-maintained computers at the beginning will be replaced by the end
  - ❖ The experiment will not “own” any of the CPU resources it uses
  - ❖ At least one migration to a new operating system will occur
  - ❖ One tape density migration will be needed
  - ❖ All students, most post-docs present at the beginning will have new jobs
    - ► Currently takes about 10 people to run all raw data and MC production
  - ❖ Number of people actively pursuing analysis will decline significantly
    - ► Number of people available for service will follow, possibly more sharply
    - ► Have a fairly clear picture of participation for first 2 years
  - ❖ Several major unanticipated operational or infrastructure support problems
    - ► Have had about one per year over past few

What needs to be preserved?

- Depends upon what we mean by “analysis”?
  - ❖ Discussed at least two possible levels of service
    - ► Full raw data reconstruction + simulation + analysis capability available now
    - ► N-tuple level analysis only: two possibilities
      - ▻ Retain ability to re-make ntuples
        - ◆ *Need to keep production output + ntuples + some (most?) reconstruction capability*
      - ▻ Only run on ntuples in existence near end of run
        - ◆ *Need to keep ntuples only*

    In either case, will likely need / want some simulation capability
  - ❖ Other possible levels of service
    - ► Simplified data structures
    - ► Generic four-vectors

  Will not consider these
    - ► Further simplification of data requires significant validation effort + new infrastructure / computing model
    - ► Would need to perform a “test analysis” under these assumptions to see if plan is viable
    - ► We don’t currently have the manpower to pursue these

Retaining Production Capability

- Offline code (reconstruction, ntupling, simulation, generation)
  - ❖ All code in frozen releases, or "tagged" in CVS repository as incremental change to frozen release
    - ► Reconstruction and detector simulation very stable over past few years
    - ► Ntupling code changes over time scale of a year
  - ❖ Build infrastructure / platforms
    - ► Currently support SL4, SL5. Will be SLx in distant future.
  - ❖ Production scripts in CVS
- Computing infrastructure (issues of funding)
  - ❖ Computing farm (Fermigrid + remote collaborating, opportunistic grid sites)
  - ❖ Data handling systems (tape access + small cache + data catalog / delivery)
  - ❖ Non-cache disk (few x 10 TB for combined production + MC)
  - ❖ Database (calibration + trigger + run / accelerator information)
  - ❖ Other associated services currently in use

Retaining Production Capability

- Documentation and knowledge retention (experiment specific)
  - ❖ Most procedures well documented
    - ► Primarily on internal web pages actively maintained by operations groups

  Have organizations that can provide for continuity in operations
    - ► Have demonstrated transfer of operational responsibility numerous times
  - ❖ Reconstruction code partly documented via: (in order of increasing fragility)
    - ► Internal documents and web pages
    - ► Comments in the code
    - ► Email archives of reconstructions groups
    - ► Personal email
    - ► People's heads

  Deep expertise becoming harder to find as original authors leave, but is not often needed
    - ► Some exposure here

Retaining analysis capability

- To retain all current capabilities
  - ❖ Need all resources required for production + additional computing infrastructure for analysis
  - ❖ Analysis code
- Ntuple analysis only
  - ❖ With capacity to re-make ntuples (private and centralized)
    - ► Requires both production operations infrastructure + significant fraction of raw data reconstruction capability
    - ► If also want simulation capability
      - ▹ Either invent a new pathway from generator to simulated ntuples, or...
      - ▹ ...need current generators + simulation + full reconstruction

    Quickly end up in the “retain all current capabilities” scenario

Retaining analysis capability

- Computing infrastructure
  - Ntuples only, no re-making (or simulation) capability
    - Computing resources
    - Data handling system (reduced size)
    - Non-cache disk (significant size)
    - Little else needed
  - Ntuples only, re-making capability
    - All the above with larger DH system
    - Ntuple production scripts
    - Most of reconstruction
    - Database

Will come back to resource requirements...

Retaining analysis capability

- Documentation and knowledge retention
  - ❖ Information relevant to analysis in
    - ► Control room logbooks
    - ► Internal web pages
    - ► Private email and logbooks
    - ► People's heads
  - ❖ Many analyses incrementally updated: analysis groups face continuity issues
    - ► Major systematic efforts within analysis groups to address this
      - ▹ Promoted adoption of common analysis frameworks, simplification of analysis infrastructure, documentation of procedures and practices, code archival requirements
      - ▹ Greatly reduced entry cost of contributing, extending analyses
      - ▹ Tremendous progress made within this area even within the past year

What resources are needed?

- Analysis demand model
  - Separate analysis into two major categories
    - “Core” analyses
    - (as defined in the Tevatron Collider Experiment Task Force Report, Dec., 2005)
    - “Other” analyses
  - Core analyses
    - Assume these are always fully staffed, so computing demand remains high
    - Continuous evolution with time
      - More complex / sophisticated algorithms (e.g., matrix element methods)
      - Better procedures or more CPU efficient algorithms
    - Need the full complement of production activities to support core analyses
  - Other analyses
    - Staffed with remaining effort
    - Demand scales with the number of people working on non-core analyses
    - May only need a small fraction of the production output
    - Can in many cases leverage MC samples produced for core analyses

Measured on-site CPU demand

Expect demand to remain approx constant through end of data taking. Then depends upon analysis needs. Short of large-scale re-processing, it should become significantly smaller

Demand scales with number of non-core analyses or active physicists. Will decline as LHC activity increases.

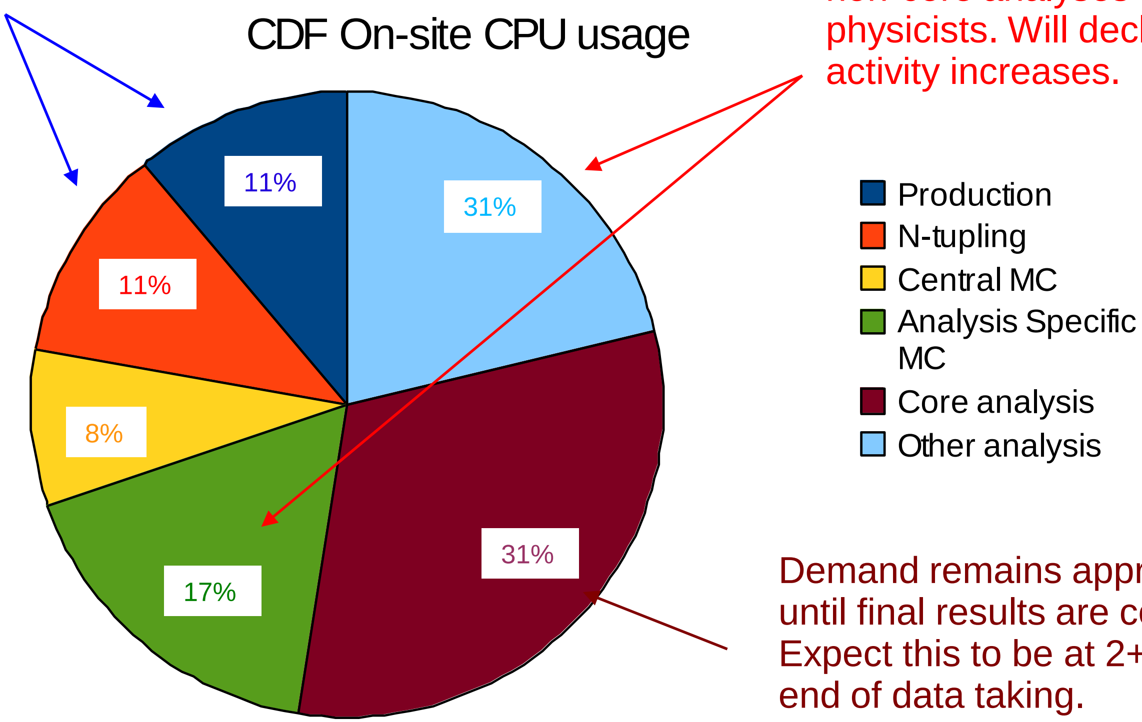

Requirements tied to size of total dataset. Expect demand to ramp down after end of data taking.

Demand remains approx constant until final results are completed. Expect this to be at 2+ years after end of data taking.

Will need about between 1/2 and 2/3 of current computing capacity to meet demand a year after the end of data taking, depending upon level of effort available and collaboration ambition.

Data preservation

- Data on tape
  - ❖ Adding about 1.5 PB/year
  - ❖ Expect to have ~9 PB by end of 2011
    - ► Approx volume by data type
      - 3,0 PB raw data
      - 4,0 PB production output
      - 1.3 PB MC data
      - 1.4 PB ntuple data
  - ❖ Technically easy but costly to preserve
  - ❖ All data is stored in root format
- Data not on tape
  - ❖ Online logbooks, internal notes and web pages, code repository, etc.
  - ❖ Mostly small volume
  - ❖ Difficulty in some cases is to identify important data, get it on tape
    - ► Physics groups now actively pursuing this to insure we get it all stored properly

Points Not Touched Upon

- Ideally for this to be successful, we would like to have an analysis expert retained – to keep systems going, modernize code as platforms change, perform test analyses to insure we get the same result as things are migrated etc. Testing all paths important
  - Important job but not very attractive career move for a physicist and difficult role for non physicist to fill
- Would like to hire a few people to look beyond the current 5 year plan – both scientist and computing professionals to be proactive in what we want to do in that phase – currently no dedicated resources
  - Lack of funding support specifically targeted for this effort means that progress is very slow and sporadic.

Summary

- Goal is to retain current capabilities for at least 5 years
  - ❖ A good physics case for this time scale
  - ❖ Initial CPU demand will be 1/2 – 2/3 current level
    - ► Period may last 2 – 3 years while core analyses are completed and published
    - ► Continue operating existing resources
    - ► Technically simple solution, but an expensive enterprise
  - ❖ Need to improve demand model to understand requirements in later years
    - ► Smaller scale, but some significant expenses accrue (HW upgrades, tape migration,....)
    - ► Additional risk due to evolution of technologies, support
  - ❖ Retaining production capability straight-forward
    - ► Some risk associated with reconstruction support
  - ❖ Retaining analysis capability is less so
    - ► Will require sustained organizational effort to ensure critical knowledge is retained
      - ▹ Many major analyses are in good shape or headed in right direction
  - ❖ Will need to secure funding to maintain computing plant, continue operations, adjust to changes in underlying infrastructure and services