Kevin McFarland
University of Rochester
for the CDF Level-3 and Offline Groups

CDF Upgrade Review
February 1, 2000

---

1. Overview of Level-3 & High-Level DAQ
2. Preparations for Acquiring Data:
   Data Flow and Integration Tests
3. Preparations for Analyzing Data:
   Reconstruction and Filtering
4. Milestones and Future Plans

Level-3 and Commissioning

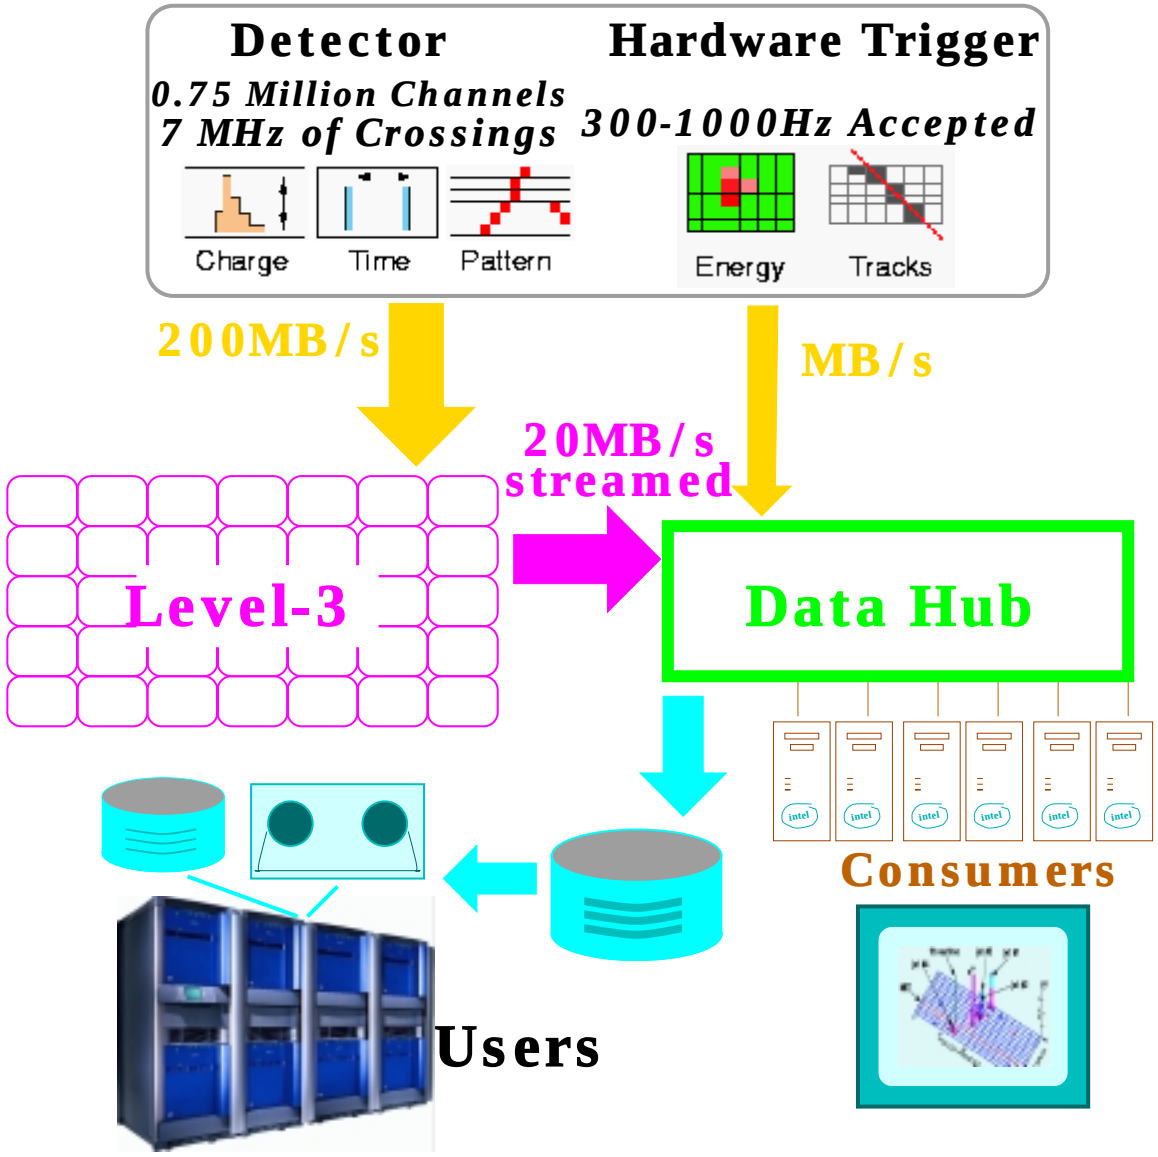

Level-3 and High-Level DAQ are. . .

. . . gateway and traffic control in the data path for commissioning

. . . where events are first assembled

. . . a point for processor-based triggering

. . . the selectors for online monitoring

Status: Assembly (a reminder)

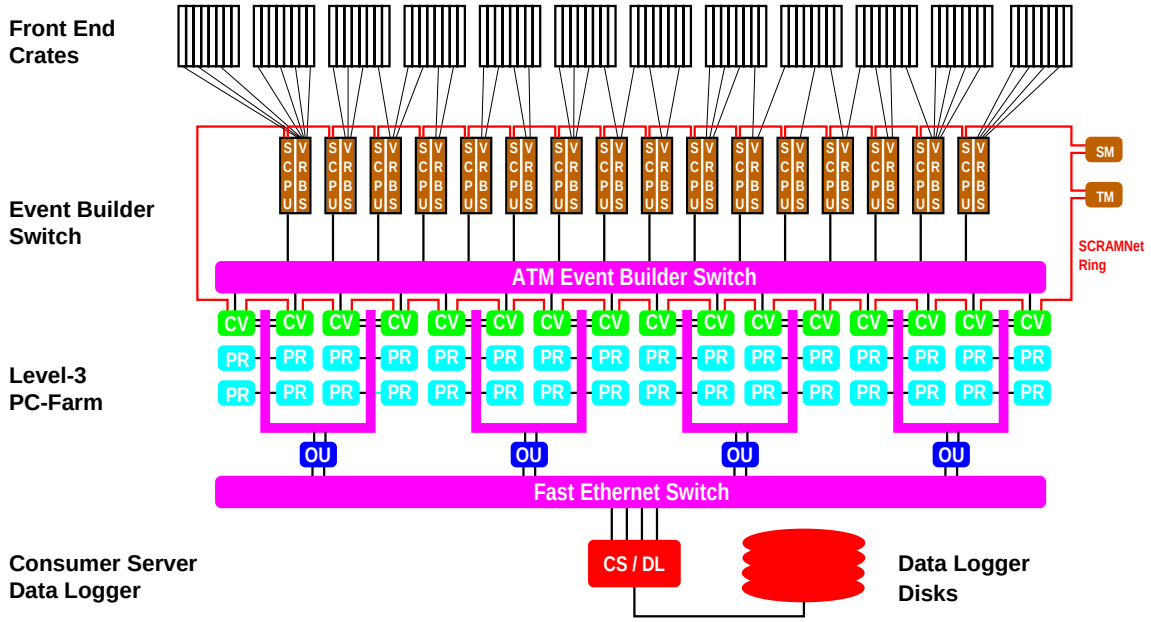

Event Builder Assembles events into Level-3 Processors

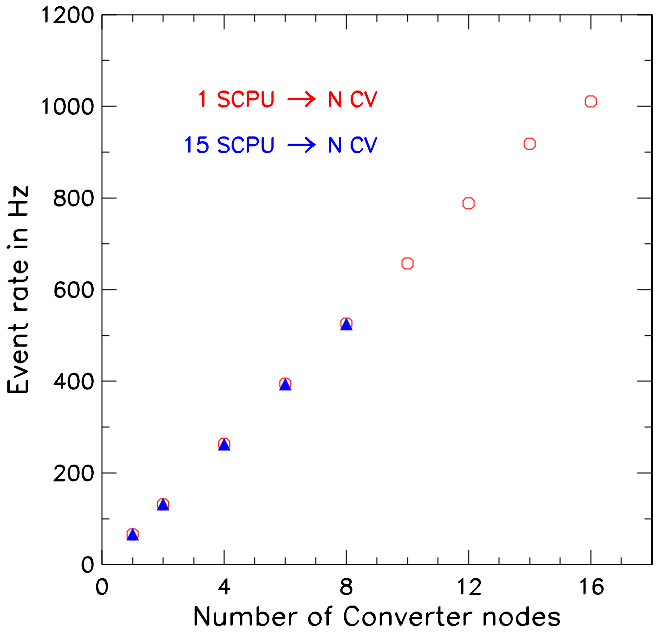

- Well on way to 200MB/s goal
- Well over design bandwidth of 75 MB/s
- Personnel: MIT, 6 FTE

Refer to Jim Patrick's talk for details

Status: Distribution

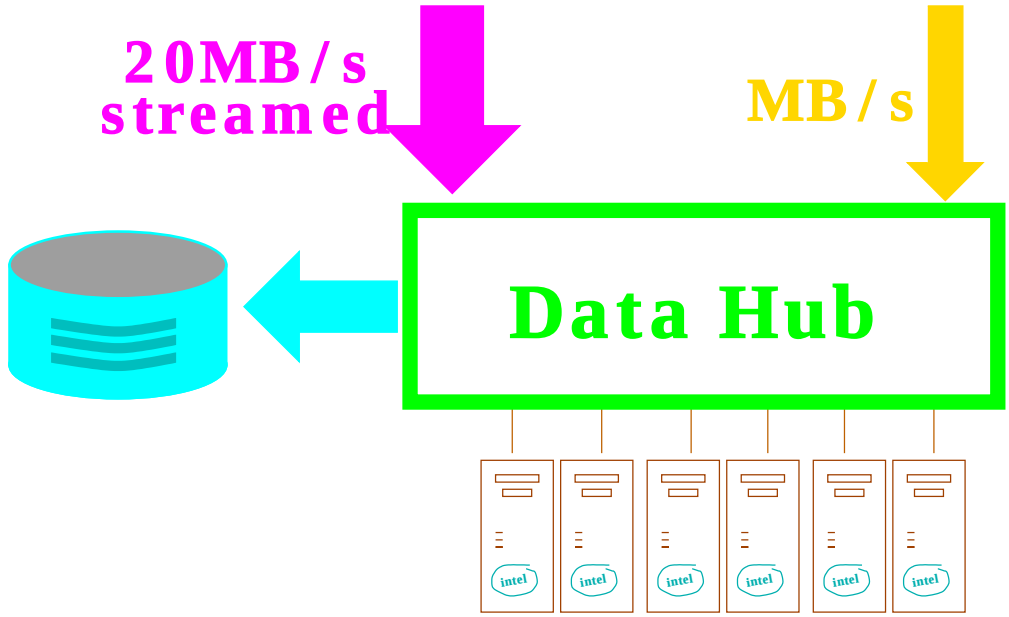

- Supports multiple data sources
  - ▷ Software Event-Builder (low rate/one stream)
  - ▷ Hardware Event-Builder (through Level-3/high rate/streams)
- Supports multiple simultaneous users **(partitions)**
- Design Throughput Achieved
  - ▷ Data$\longrightarrow$Disk (8 sources) 32 MB/s
  - ▷ 5 MB/s to Consumers feasible; needs more realistic tests
  - ▷ Support for all offline data formats (TRYBOS, ROOT)
  - ▷ Supports streamed output as defined by Level-3
- Event Distribution to Monitors done
  - ▷ Each monitor currently regulates own rate
  - ▷ Prioritization/selection scheme under active development

Status: Distribution (cont'd)

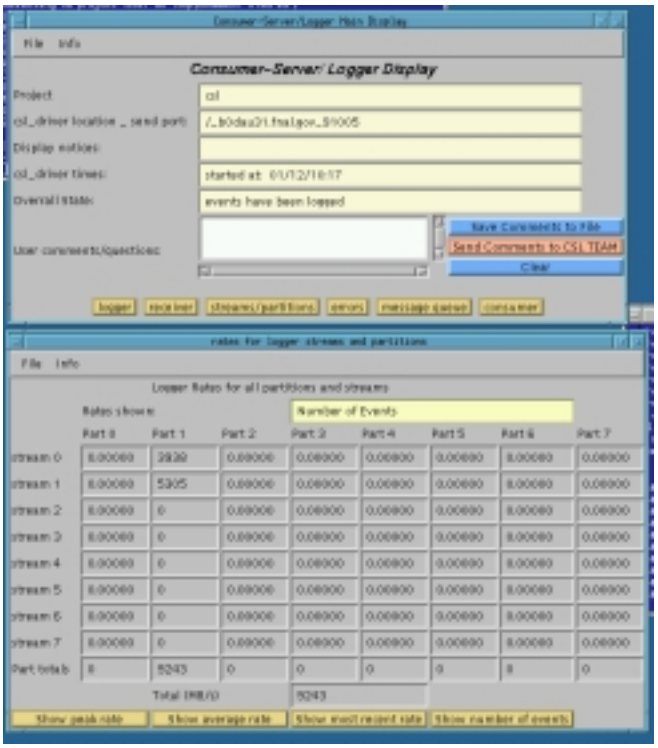

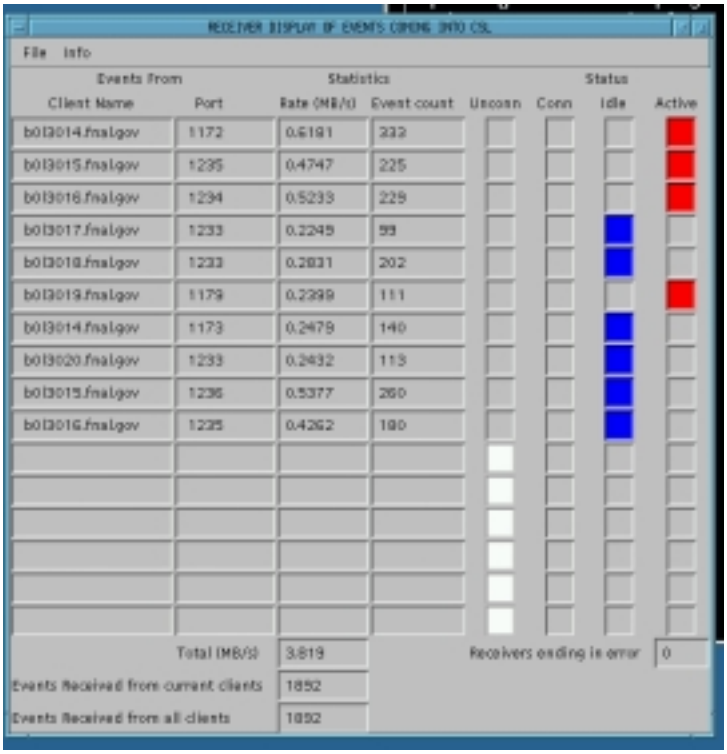

- Operations
  - ▷ Currently supporting 24×7 operation (with pager)
  - ▷ No automated start-up/shut-down
  - ▷ Monitoring available
- Data Hub Personnel
  - ▷ University of Rochester, 1.5 FTE
  - ▷ University of Tsukuba, 1 FTE
  - ▷ Support/development currently 50–50

Status: Analysis, Online

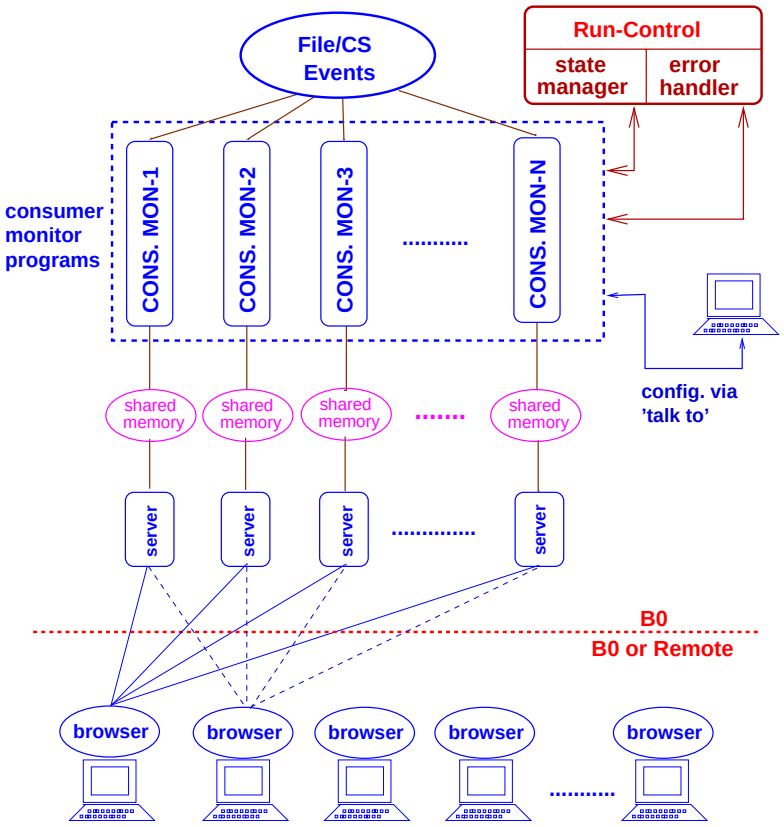

- Multiple Consumer PCs
  - ▷ Different analyses and events
- ROOT-based shared memory/Common Browser
- Framework Status
  - ▷ Uses offline framework for data access, reconstruction *(more later)*
  - ▷ Memory/browser operational, under active development
- Personnel
  - ▷ $\sim$ 2 FTE (Karlsruhe, FNAL, Waseda, LBL)

Status: Analysis, Online (cont'd)

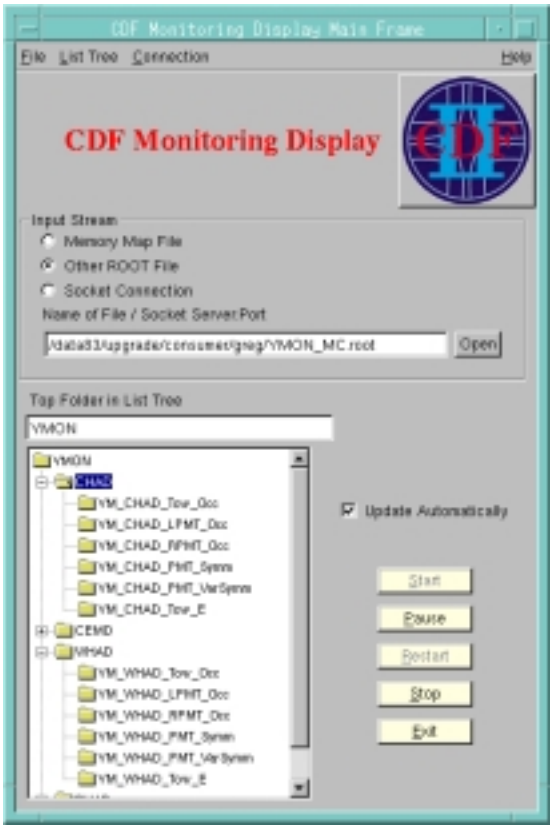

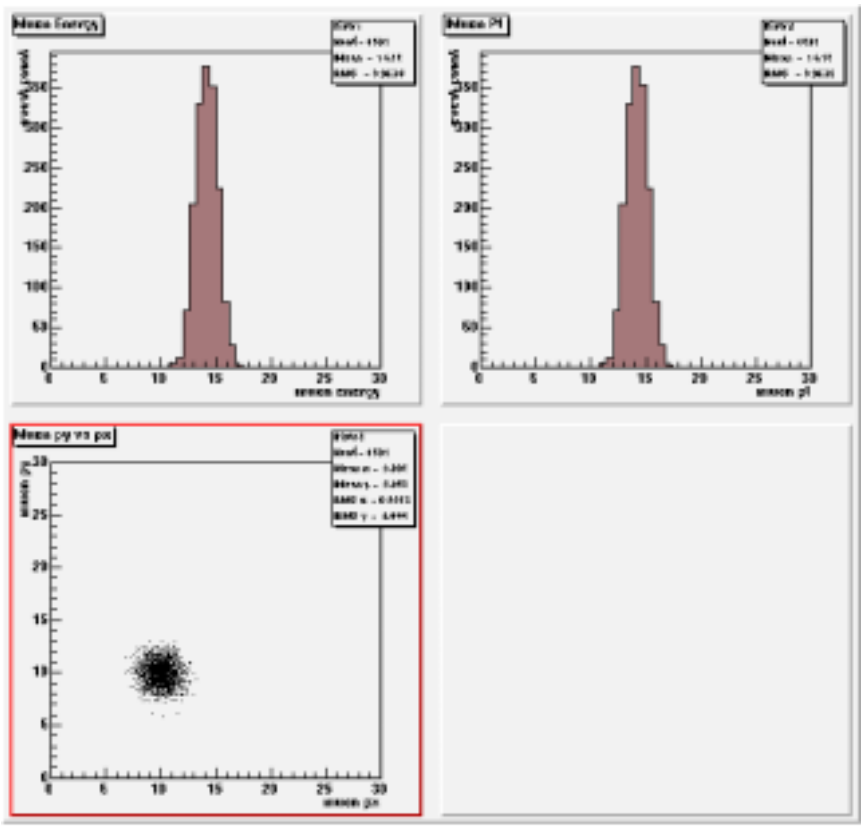

Consumer	Contributing Groups	Status
YMON	ANL, LBL	
TRIGMON	UCLA	
XMON	Michigan, Rochester	
LUMMON	Florida	
PHYSMON	Chicago	
SVX	Yale	
SVT	Roma I	
L3MON	Rochester, Oxford	
STAGE0/POST0	Harvard	
ADC-Calib	Texas Tech, FNAL	
TDC-Calib	Duke, LBL	
SVX-Calib	Glasgow	

Coding Started — $> 50\%$ of First Version Coded

tatus: Analysis, Online (cont'd)

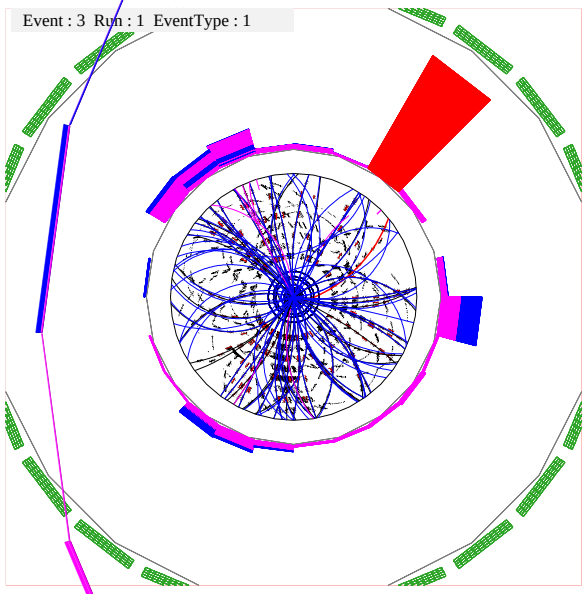

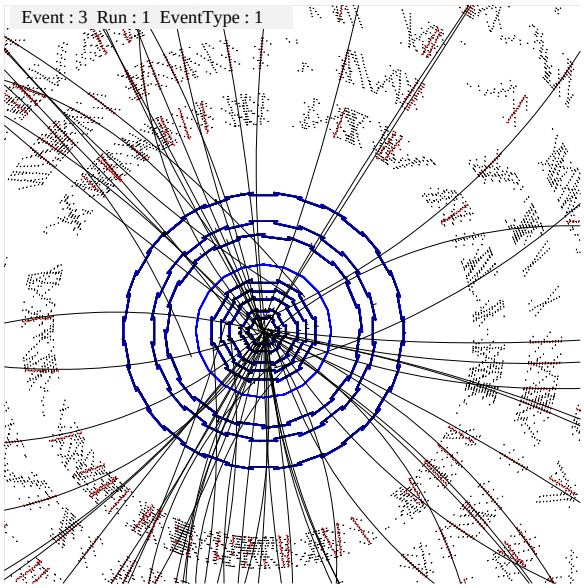

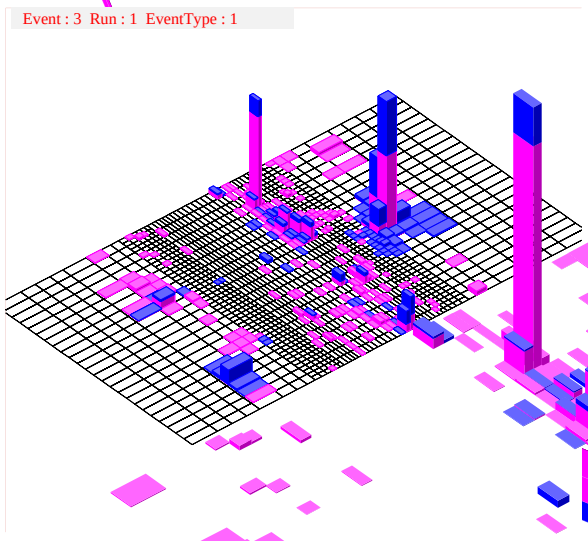

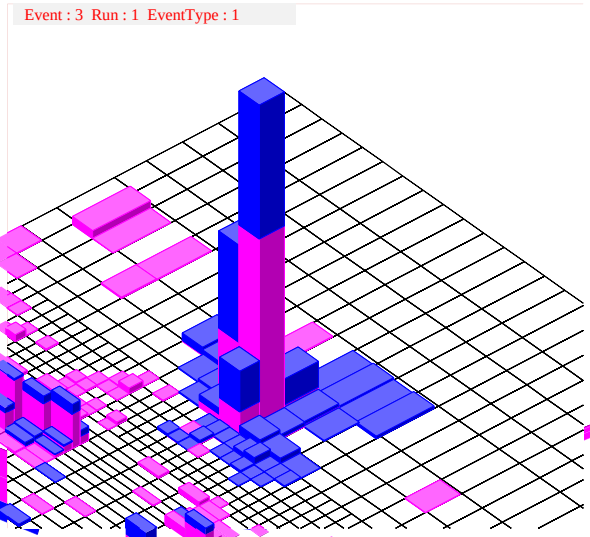

- Interactive Event Display
  - ▷ Integrated with data hub
  - ▷ Still needed: event selection, user experience
- Personnel
  - ▷ ITEP, 1 FTE

Successful Tests: Distribution & Consumers

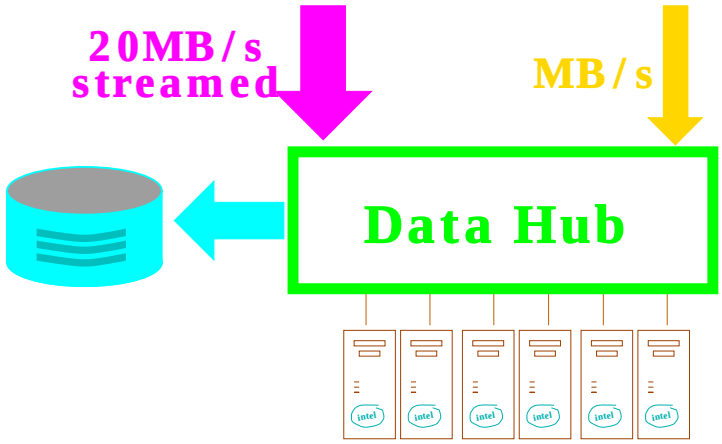

- Calorimeter Calibration Consumer
  Front-End $\longrightarrow$ Data Hub $\longrightarrow$ Consumer $\longrightarrow$ Database
  - ▷ Routinely operated
  - ▷ Results being analyzed for gain stability studies
- Event Display
  Front-End $\longrightarrow$ Level-3 $\longrightarrow$ Data Hub $\longrightarrow$ Consumer
  - ▷ Calorimeter LEGO from detector
  - ▷ Calorimeter and COT for MC events from Level-3
- Silicon Calibration Consumer
  Front-End $\longrightarrow$ Level-3 $\longrightarrow$ Data Hub $\longrightarrow$ Consumer $\longrightarrow$ Database $\longrightarrow$ Front-End (update)
  - ▷ First success mid-January
- COT Calibration Consumer
  TDC Test Crate $\longrightarrow$ Data Hub $\longrightarrow$ Consumer $\longrightarrow$ Database
  - ▷ First success mid-January

Successful Tests: Distribution & Consumers

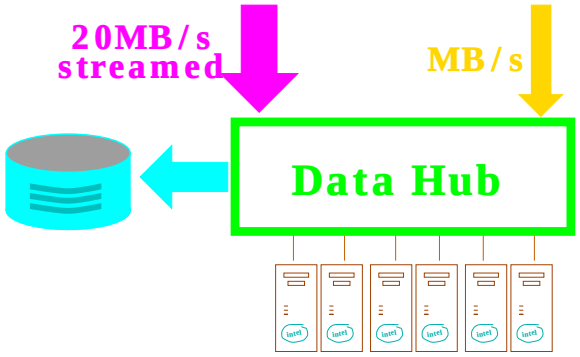

- Operation of Calorimeter Wedge on Cosmics
  Front-End (triggered!) $\longrightarrow$ Data Hub $\longrightarrow$ Event-Display

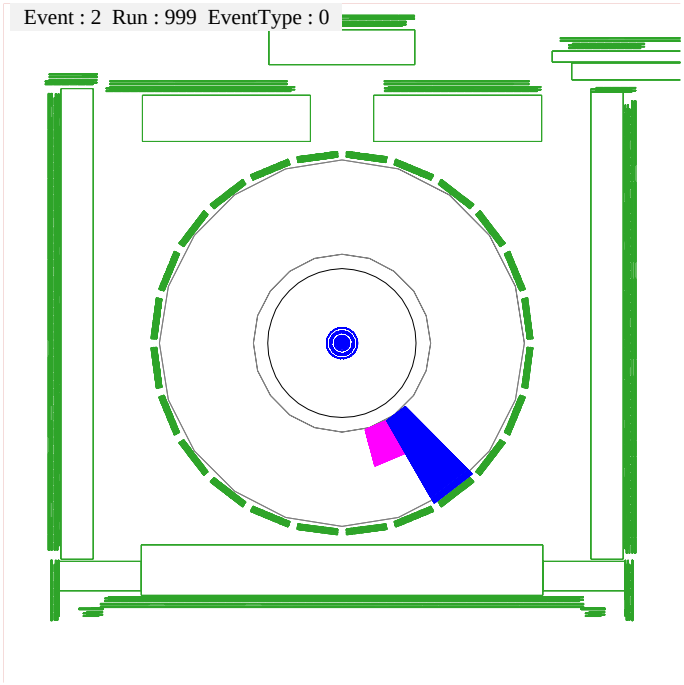

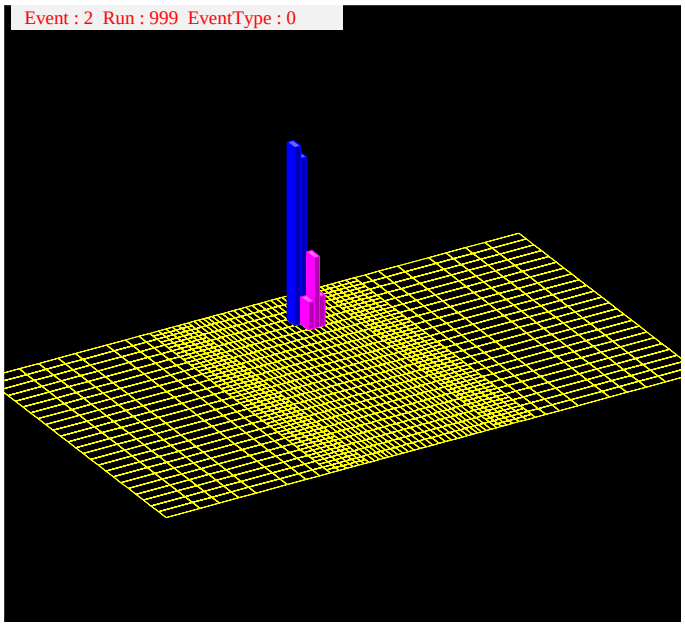

Summary of Dataflow Tests:

- Data from front-end to online applications $\surd$
- In progress: operational ease, multi-system rate

Status: Selection and Streaming

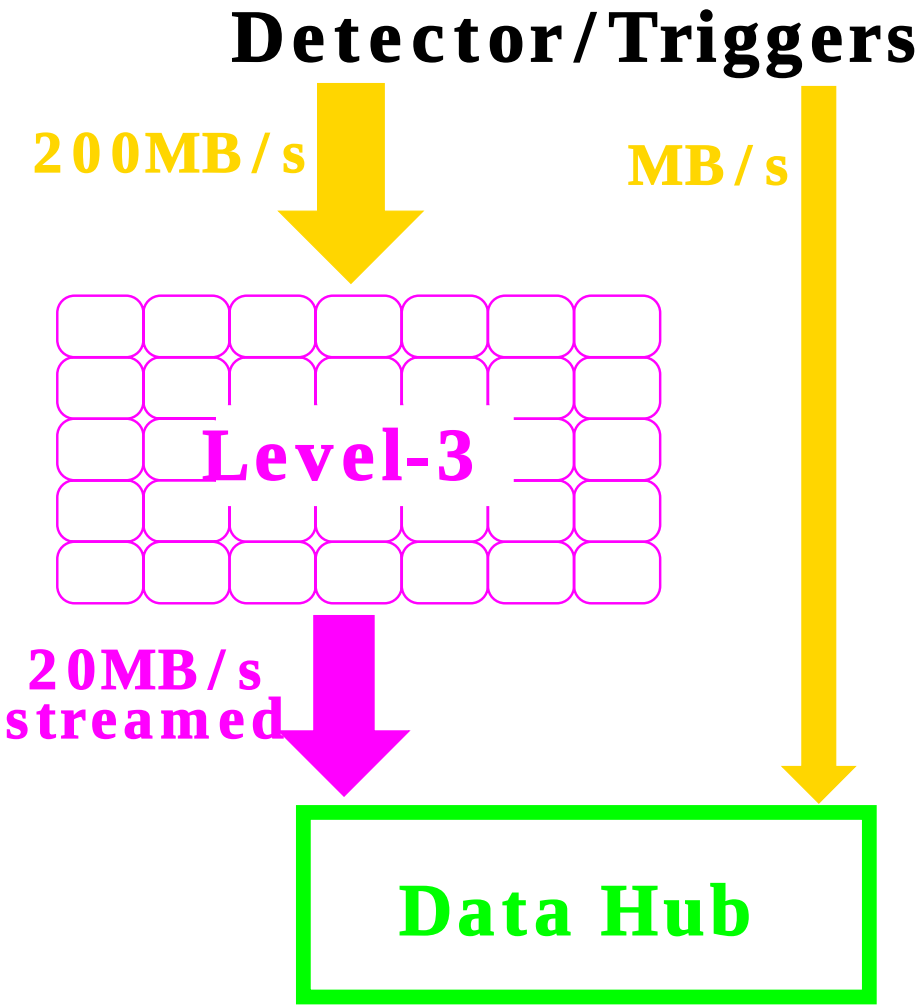

Goals:

. . . Reduced dataset to 75 Hz logging limit (20 MB/s)

. . . Prioritization and Sort remaining data in streams
(drink from the firehose)

Reconstruction infrastructure later...

Streaming from Level-3

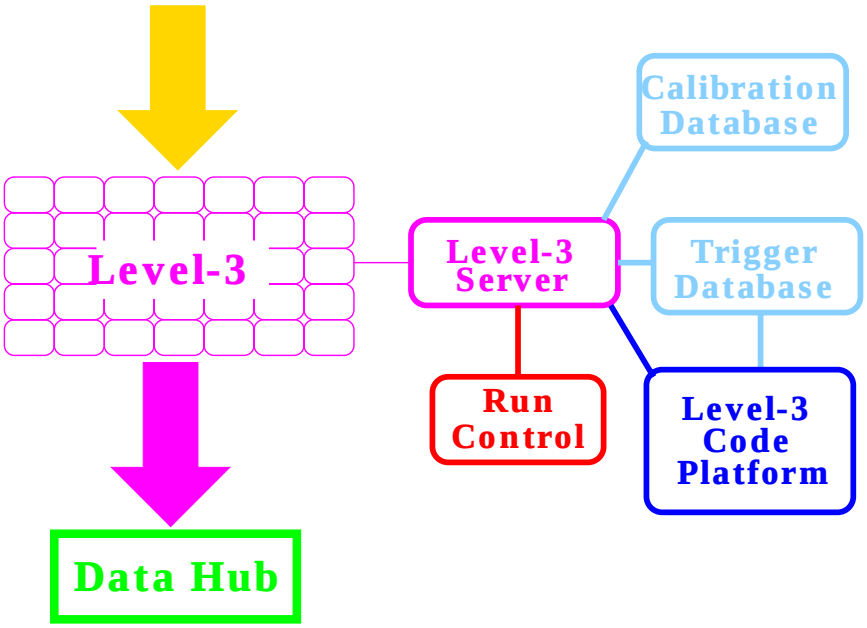

- Run start/stop
- Calibration snapshot
- Triggers and datastreams
- (Filters and reconstruction)

Trigger Database and Streaming are operational

- Tables stored in Oracle
- Pre-run
  - ▷ Level-3 server queries database
  - ▷ Information distributed into Level-3
- During run:
  - ▷ Selections determine data streams, based on trigger information
  - ▷ Data Hub takes information from Level-3 then distributes events accordingly

Building Triggers

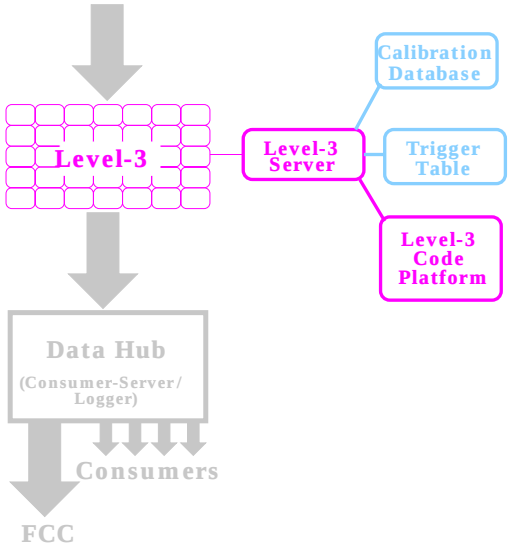

Trigger Database

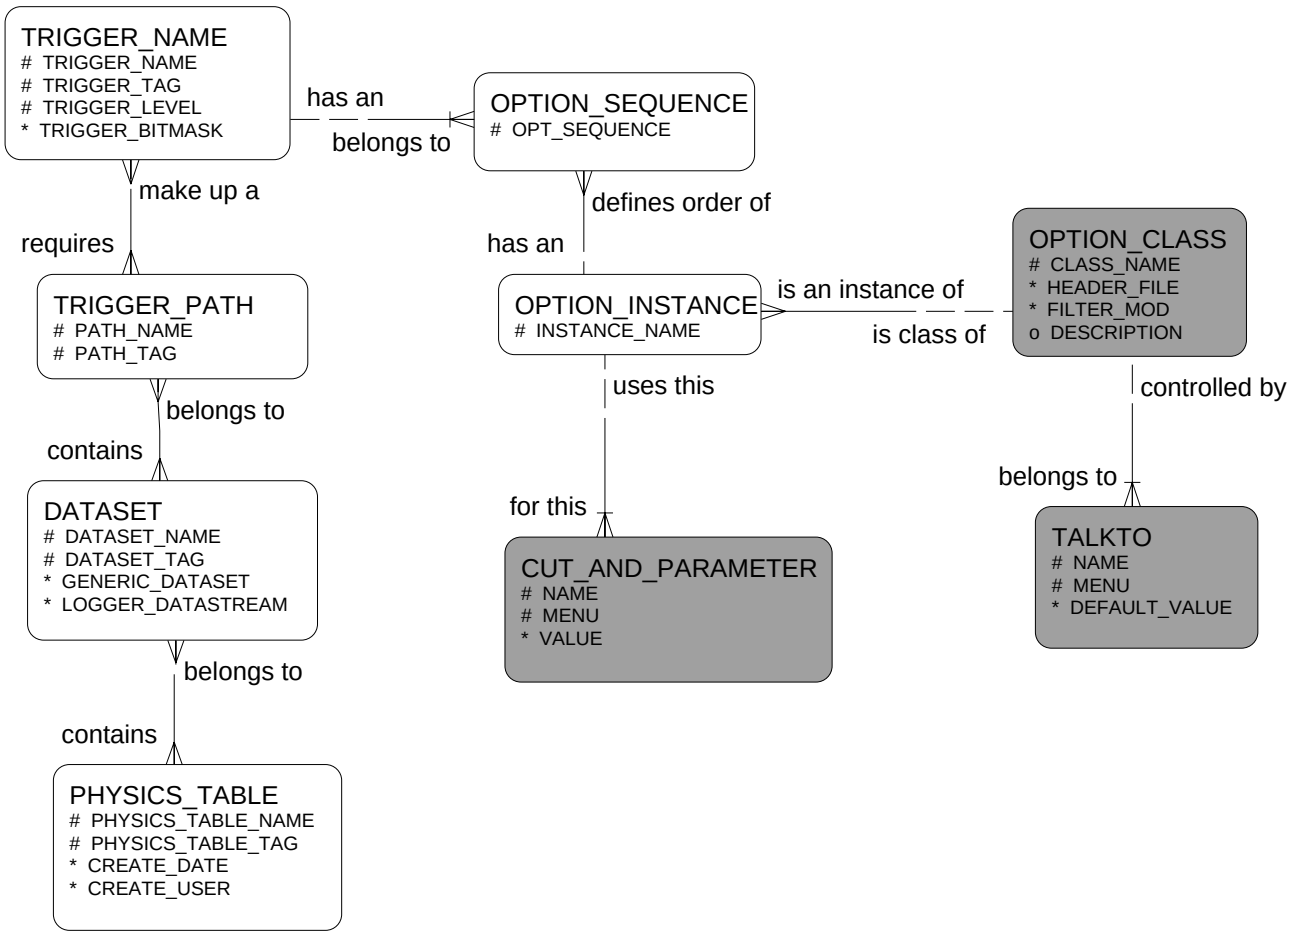

Status: Offline Data Handling

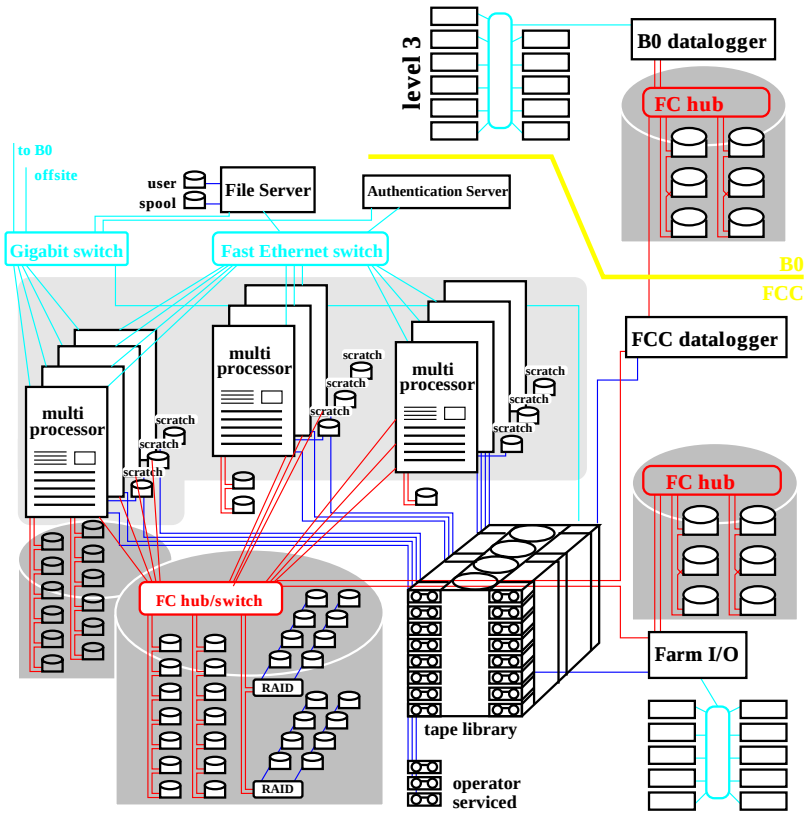

- Data logging from B0 to tape
- PC Farm for production
- Analysis SMPs
  - ▷ Multi-architecture
  - ▷ Large pools of shared disk
  - ▷ Resource management for tape/disk access
- *Data path seeded by Level-3 and production decisions*

Mock Data Challenge I

Hardware Status:

- Small Level-3, Production Farms
- Data Hub
- Offline Data Handling and SMP

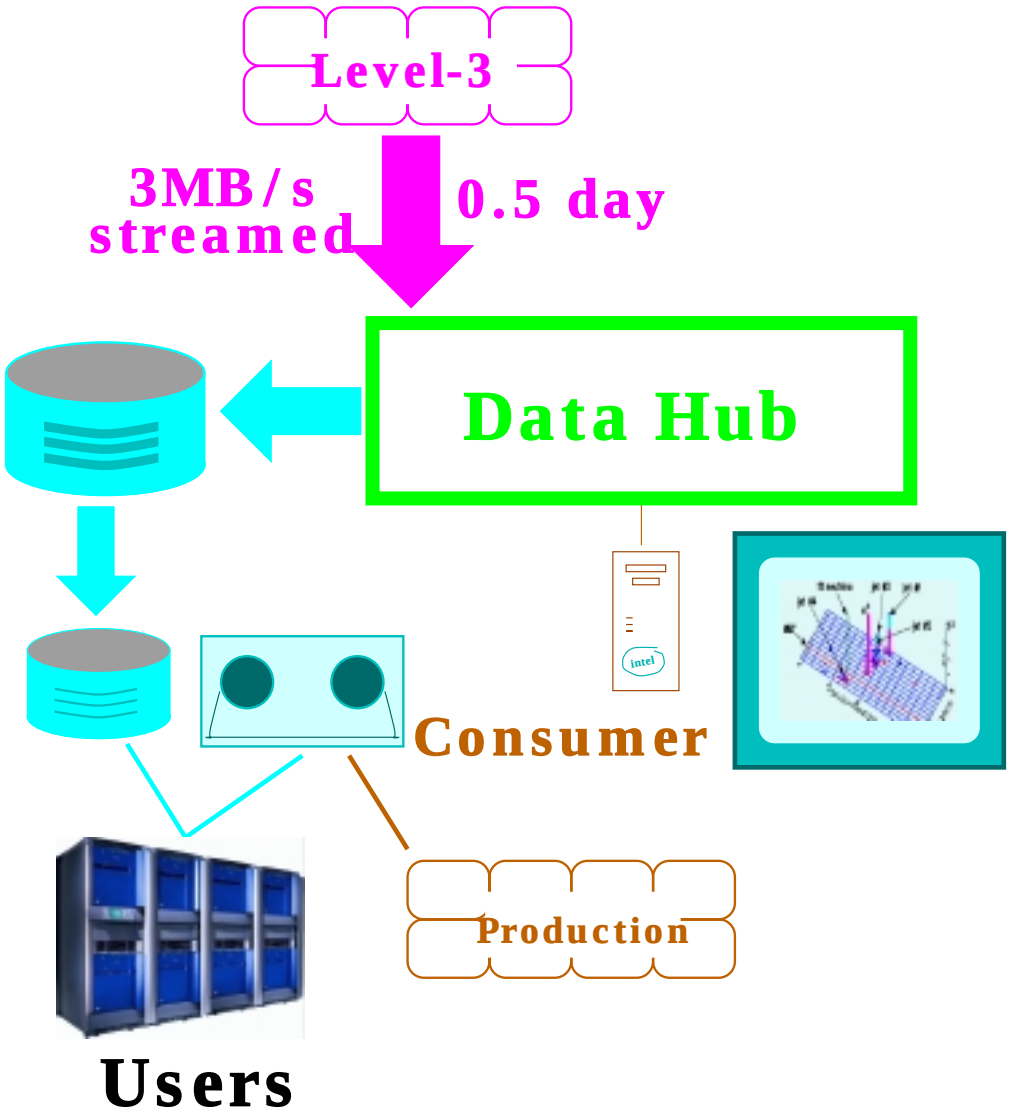

Mock Data Challenge (cont'd)

Contemplating the Mock Data Challenge

Goals:

- $\sim$ 100 GB (Monte Carlo) data processed
- Low rate OK (few MB/s in each system)
- Reconstruction and filtering through high-level reconstruction
- Many ($\approx$ 10) triggers into two data streams
- Connectivity from Level-3 through user analysis of data

MDC Streaming Status

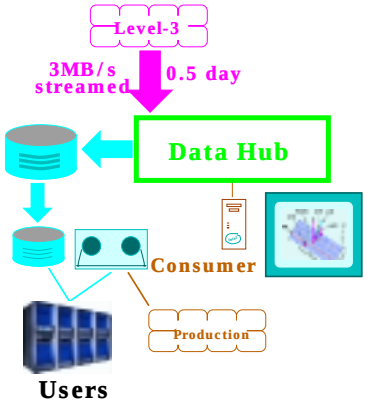

- Progress:
  - ▷ Data processed through Level-3 into two production streams
  - ▷ Archiving and production successful
- Remaining glitches
  - ▷ Data hub/logger share a dual-ported FC disk
    controller problems require logger to send data over alternate (slow) data path
  - ▷ Haven't met user analysis milestone yet
    expected mid-February
- Other Issues:
  - ▷ Trigger database works well
    . . . but integration with hardware triggers expected to be non-trivial
    (non-expert trigger entry a particular concern)
  - ▷ Questions of rate not addressed
  - ▷ Production needs to use trigger database as input
    (expect to share Level-3 mechanism)

Status: Reconstruction

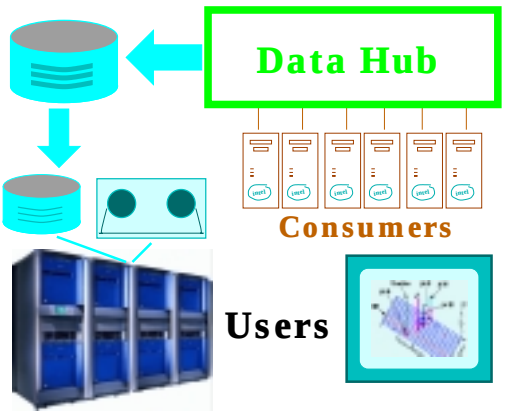

- *Level-3*, *Consumers (online)* and *Offline*
  - ▷ Common genesis in Offline framework
- Level-3
  - ▷ Low- (tracks, clusters) and High-Level (leptons, vertices) reconstruction
  - ▷ Feed datastreams based on these results
  - ▷ *Level-3 analyzes* $10\times$ *the data on tape!*
- Consumers
  - ▷ Raw Data, Low-level & High-Level monitoring
  - ▷ Trigger monitoring & validation
  - ▷ Event display (reconstructed quantities!)
  - ▷ Calibrations
- Offline
  - ▷ Code (monitoring, calibration) development
  - ▷ Eventually, production

Algorithmic Framework for Level-3

Demands on Level-3 CPU

- Lepton/photon verification
  (seeded by hardware triggers)
- Heavy flavor tagging
  (often seeded by soft leptons, tracks, SVT)
- Multi-body reconstruction
  (often of elements from hardware triggers)

$\Longrightarrow$ Level-3 often driven by hardware trigger seeds

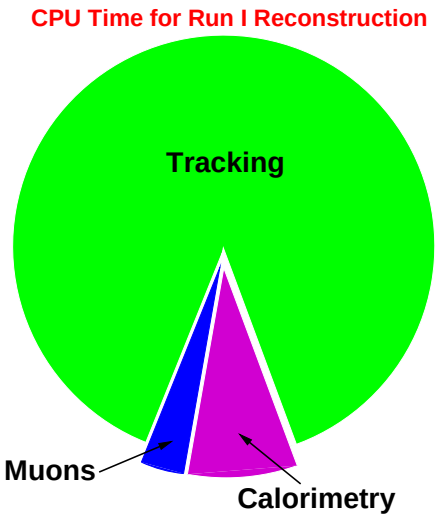

Regional Reconstruction

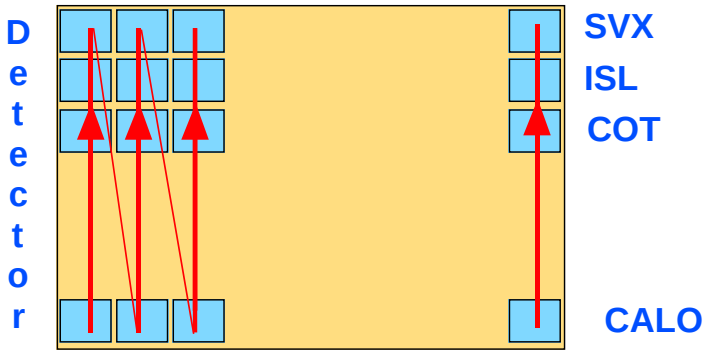

SVX/ISL and COT clear candidates for reconstruction with Seeded or Regional algorithms

Development focussed in areas new to Run II Level-3

- Si reconstruction (heavy flavor)
- Not, COT and calorimetry (Run I)

Operating Level-3

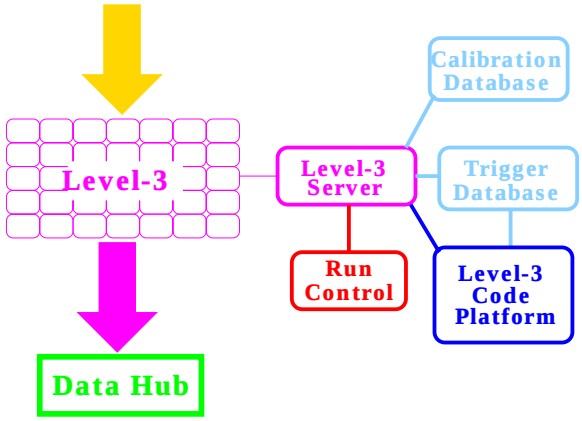

1. Pre-run

- Build Executable
  - ▷ Access trigger database to determine desired Level-3 analysis modules. filters and trigger prerequisites
  - ▷ Build (automated) and validate reconstruction
  - ▷ Export of executable and necessary shared libraries (no network access)
- Calibration Database Information
  - ▷ Tag calibration set/validate
  - ▷ Export to local (flat) database (no network access)

2. Run-Time

- Trigger/datastream results needed by Data Hub
- Filter Decision Monitoring (`XMON`-level)
  - ▷ Pass/fail probabilities by trigger and filter module
  - ▷ Resources by trigger and filter module

Operating Level-3 (cont'd)

- Event Representation Conversion

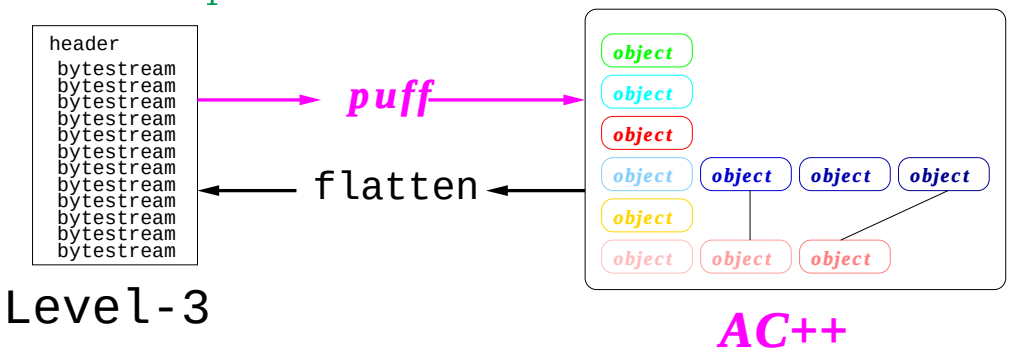

  - ▷ Data assembled from VRBs in YBOS-like mini-banks
  - ▷ TRYBOS (Banks) $\longrightarrow$ EDM-2 (Objects/ROOT)

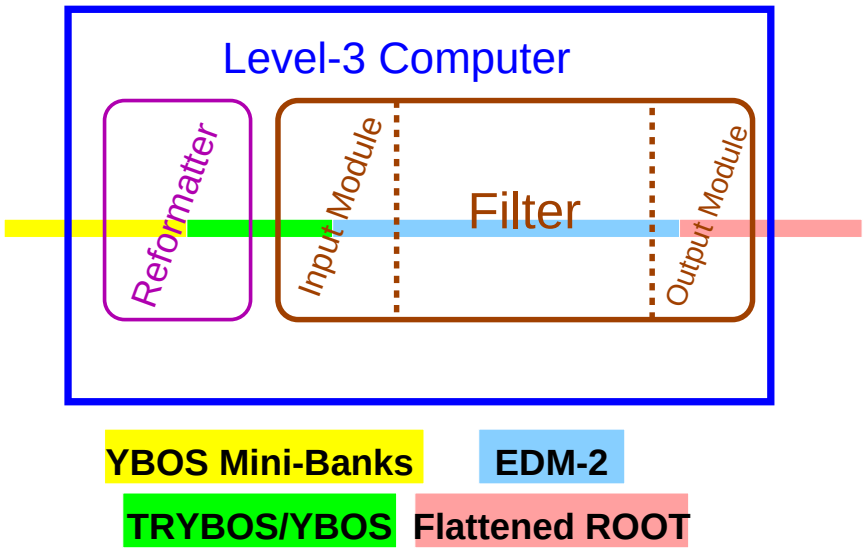

  - ▷ Output of event is "flattened" ROOT, similar to the ROOT sequential file format
    (TRYBOS output capability preserved)

- Personnel for Level-3-specific Infrastructure/Reconstruction
  Rochester (2 FTE), Oxford (2.5 FTE)

Reconstruction (Offline Project)

1. Infrastructure intact

- Recent major upgrade to ROOT-based OO data model complete
- Simulation/reconstruction geometry integration ongoing

2. Low-level (α-version) complete

- Silicon & COT Tracking, Calorimetry
  in advanced study stage
- EM reconstruction, muons functional

3. High-level reconstruction

- `CdfTrack`, lepton, jet, $\not\!\!E_T$ objects all exist
- First filter and analysis code using these objects is complete

4. Integration/maintenance efforts

- Calibration code and database interface
- Full offline production executable assembled and validated nightly

5. Level-3: Regional COT/SVX/ISL tracking implemented

- Global COT, Si use up to ~ 1, ~ 10 CPU-sec, respectively ($t\bar{t}$ MC)
- Regional COT ~ 50 msec to track seeds from hardware track trigger (XFT)
- Regional Si
  ~ 25 msec per 3D COT track seed
- CPU savings varies by seed and implementation
- Development of examples ongoing

Calibration Database

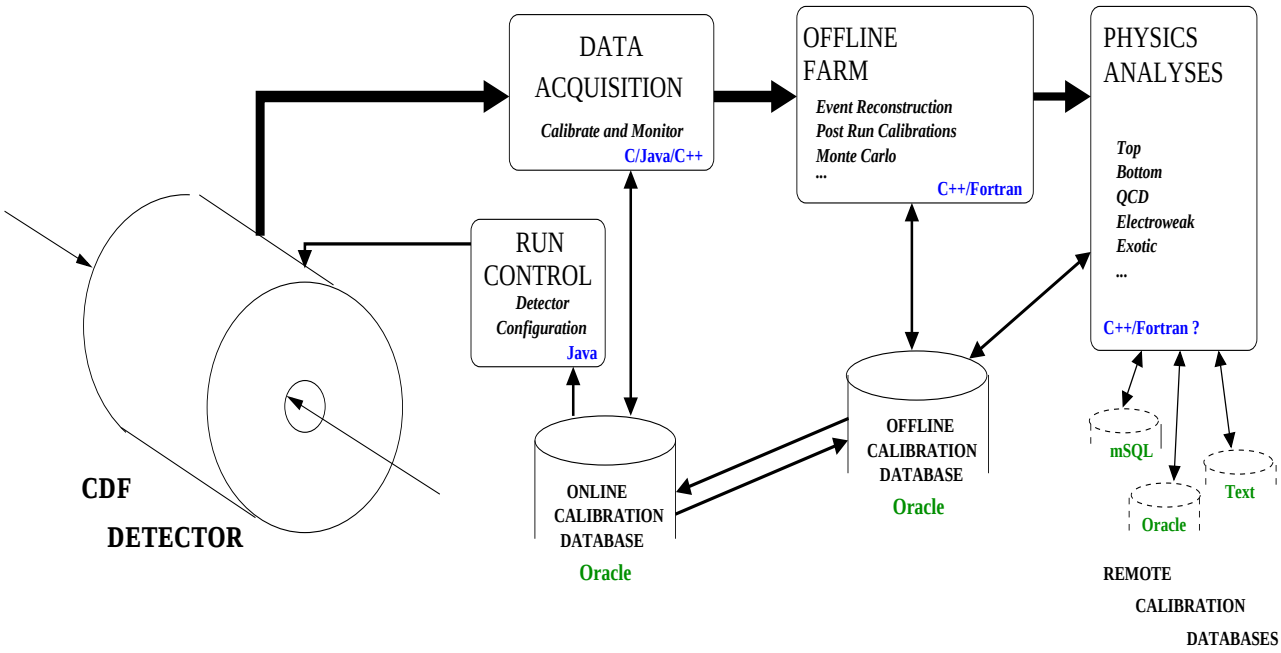

Front End Broker supplies validated constants to the front end using `Java` Calibration API

CalibConsumers write data from detector front-end to database, either with or without an analysis step. `C++` calibration API

DBANA is a ROOT-based GUI for browsing data (currently only calorimeter) using `C++` calibration API

Offline Database maintained by Oracle replication between servers in B0 and FCC.

Data Sorting/Merging determines and tags good calibrations; also can merge partial calibrations to make a complete channel set

Tracking

COT strategy well-developed:

- Find **axial tracks** from super-layer segments
- Attach stereo segments when simple
- Find z-vertex
- Revisit stereo segments with z constraint

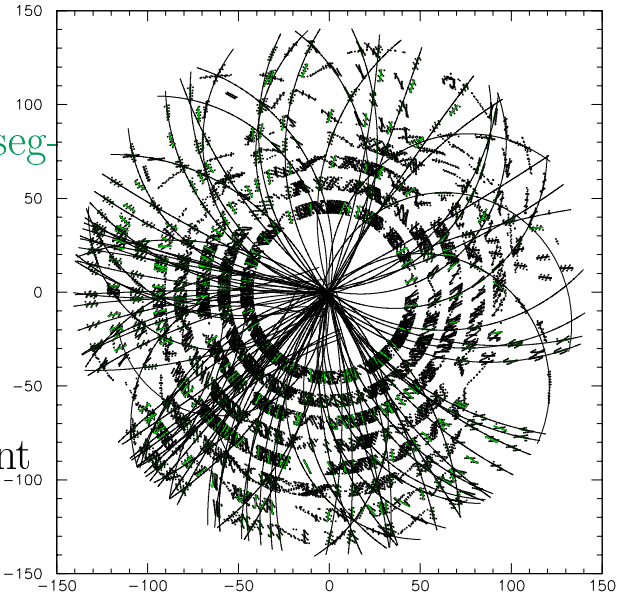

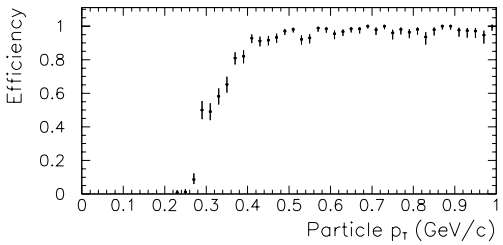

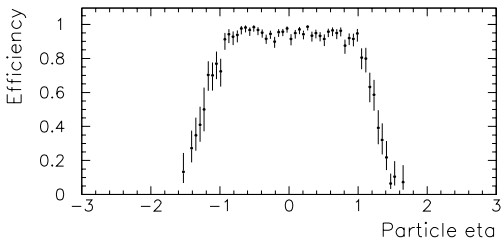

Efficiency/MC Studies Ongoing

- 96% efficient
  but "Wrong" track ($> 5\sigma$ in fit parameter) found $\approx 8\%$ of time
- *Complementary algorithm* being investigated
- Investigating efficiency for 132 ns operation

Cosmic Event

- In full-length prototype chamber
- Developing algorithms efficient on cosmics in full COT

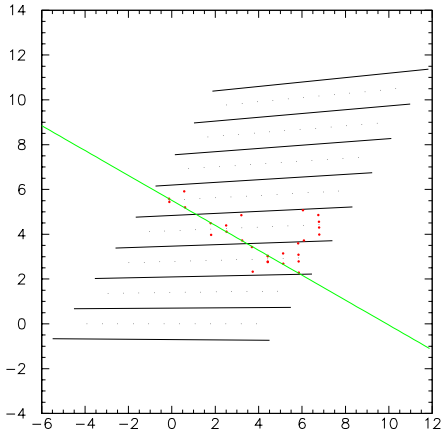

Tracking (cont'd)

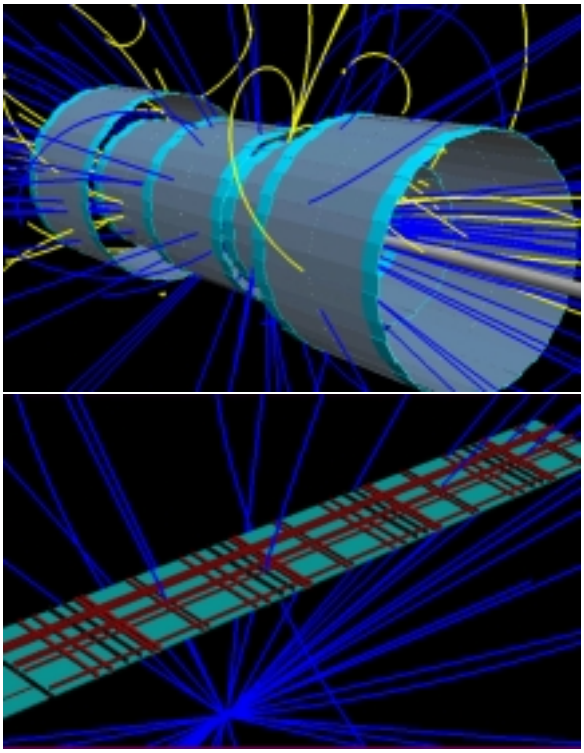

SVX/ISL

- Many layer, double-sided Si
- Substantial stand-alone capabilities

Efficiency

- $t\bar{t}$ MC
- Little fall-off out to
  $|\eta| < 1.8$ and
  $p_t > 400$ MeV

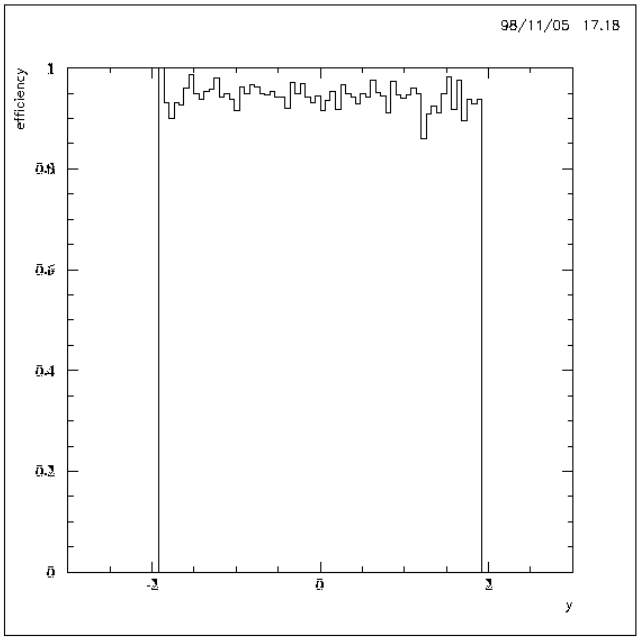

Tracking (cont'd)

Track Finding

- Outside-In (from COT)
- Stand-alone
- COT only ($K_S \to \pi\pi$ here)

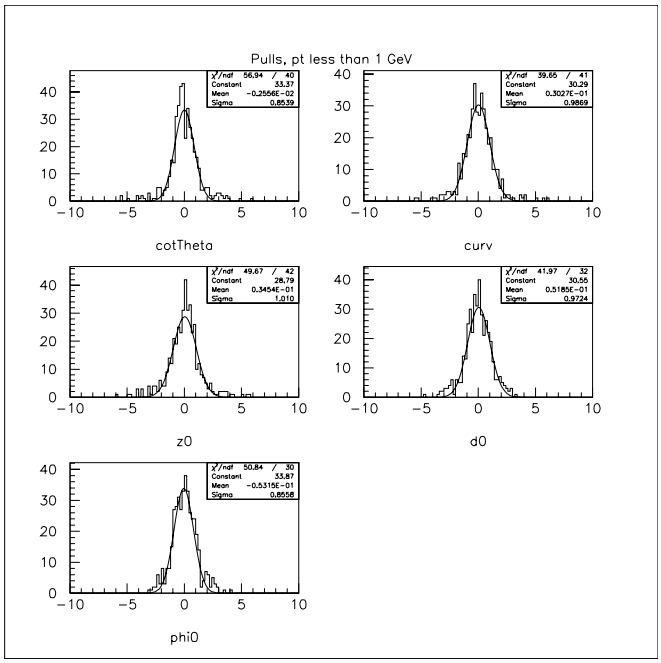

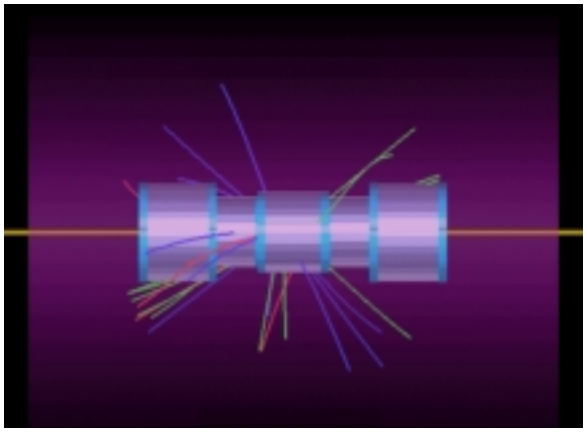

Track Fitting

- Multiple scattering in fit (low p)
- Track parameters well understood

Tracking (cont'd)

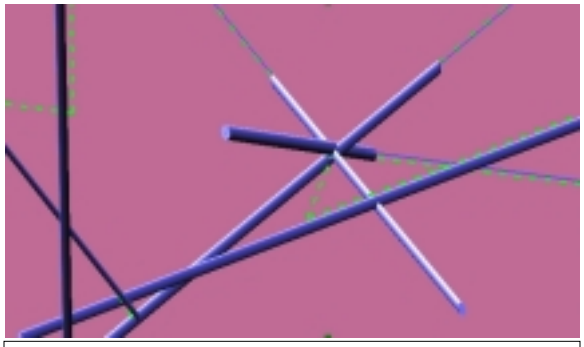

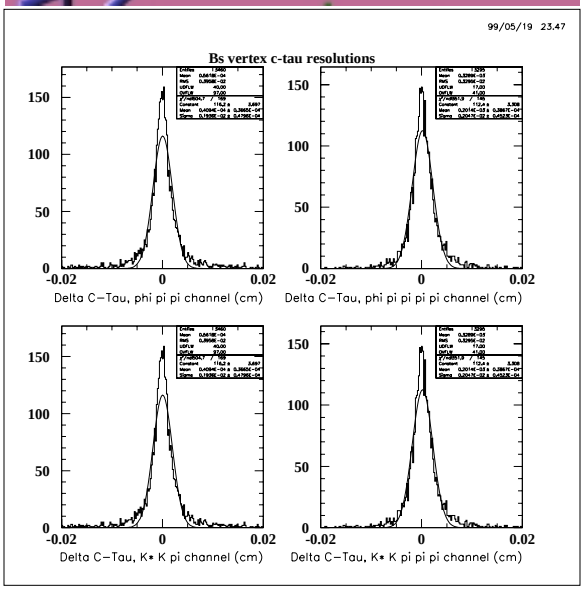

Vertex Resolution Studies

- Reconstruction useful for physics studies
- $c\tau$ resolution for hadronic B_s reconstruction
  (no underlying event or fragmentation)

B_s mixing in hadronic modes

- Resolve several periods of oscillation at $x_s = 15$
- Cheats: perfect tagging, fragmentation, underlying event, etc.
- *Still*, impressive test of physics capability

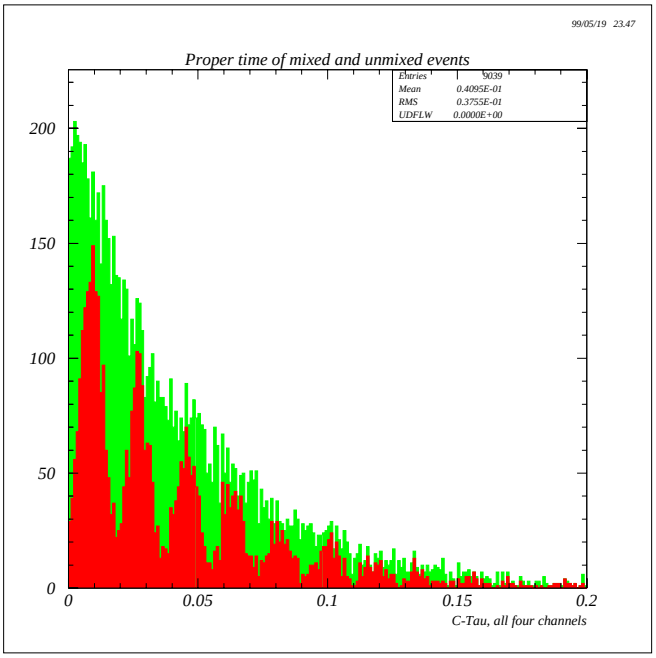

Regional Tracking

Regional Silicon Tracking

- Example: Electrons in the plug calorimeter
- Out of the range of the XFT (COT-tracking trigger)
- Can use calorimeter seed to search for forward track in SVX/ISL

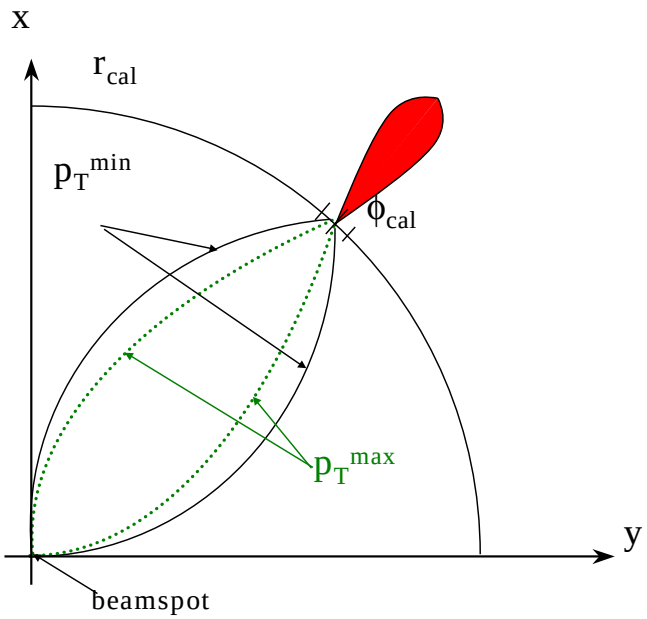

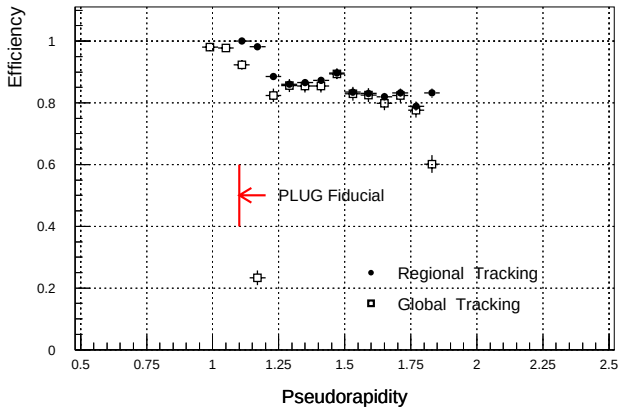

Efficiency

- Same as global
- CPU: $\sim$ 100 ms per seed

COT Regional Tracking

- Seed from XFT
- Very fast (50 ms/event); applicable for $\ell + X$ triggers
- Tested on Run I $\Psi \to \mu\mu$ (simulated XFT)

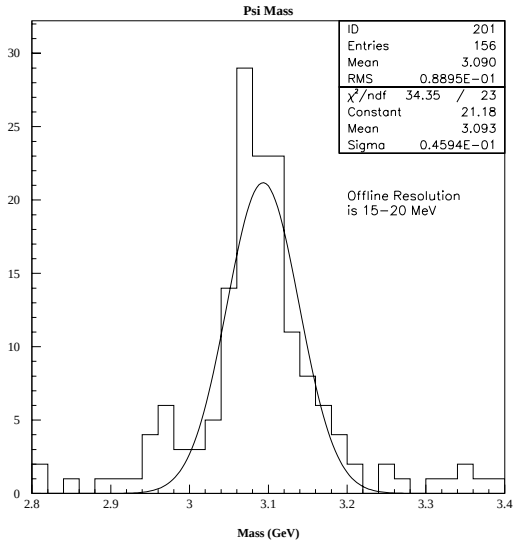

Trigger Bandwidth and CPU Usage

What if Level-2 saturates available Bandwidth ($\sim$ 200 MB/s)?...

- Evaluate high luminosity with aggressive B-physics triggers
- with minimal required Level-3 reconstruction

Specification	Design	 Saturated Scenario
Level-2 Accept	300 Hz	800 Hz
Level-3 Accept	30 Hz	75 Hz
CPU/event (PII/400MHz/512kB-L2)	0.44 sec	0.5 sec

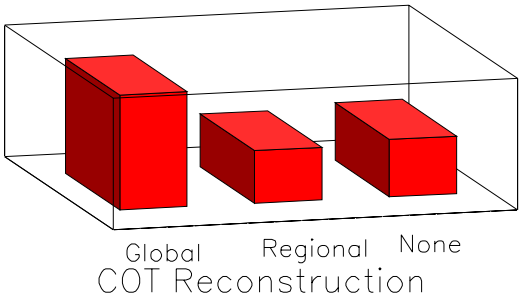

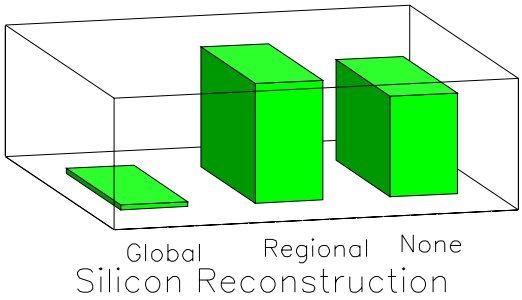

- Required rejection rate is $\sim$ 10
- CPU/event in line with expectations
  (because of regional tracking use)
  $\sim$ 700 CPU·sec/sec for March 2001
- Level-2-accepts dramatically oversubscribed...
  ...but still achievable on Level-3 side

MDC Filter Status

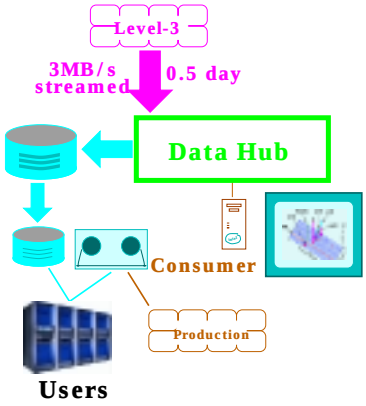

- Progress
  - ▷ 10^5 Monte Carlo events processed through Level-3 into production
  - ▷ No reconstruction crashes
  - ▷ Reconstruction feeds data streaming/splitting
- Remaining glitches
  - ▷ Haven't met user analysis milestone yet
    expected mid-Feburary
- Other Issues:
  - ▷ Calibration database access unsuccessful
    (requires some infrastructure updates to cope with changes in Oracle database schema)
  - ▷ MDC Experience in bug-hunting indicates that $\sim$ 1.5 FTE is a reasonable level of effort when filter updates are needed
    (Rochester and Oxford groups have commitments)
    $\rightarrow$ Better debugging tools under Linux needed
    CDF/D0/CD common concern
    (evaluation of new product, TotalView, underway)

Mock Data Challenge II

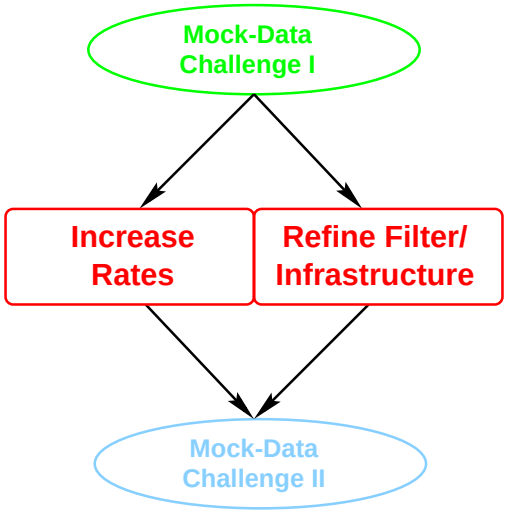

MDC I $\longrightarrow$ II

- Rate: subsystem and integrated
- Filter: more triggers, improved reconstruction
- Infrastructure: ROOT files out, complete EDM-2 integration

Schedule

- Data Hub and Level-3 farm output already at design bandwidth
  (improve consumer bandwidth)
- Production farm expanded and operational by July
  (CPU bandwidth)
- ROOT file format: have loose commitment from ROOT team
  (will proceed without this if necessary)
- Data Handling full rate tests from Level-3 to archive between mid-April and mid-May
- Full system with filter, mid-May through July

Goals through Commissioning Run

- Operation fully automated
  - ▷ Data Hub nearly there (runs from Run Control simulation)
  - ▷ Level-3 expected by March
  - ▷ FCC Data Handling:
    - → March 1: Archiving of online data and readback functional (one-time requests)
    - → April 1: Archiving of online data and subsequent readback moves to 24×7 support (few GB/day maximum)
- Increase Bandwidth to 20 MB/s
  - ▷ *Level-3 down achieved as part of MDC-II*
  - ▷ Front-End⟶Level-3⟶Data Hub⟶Archive tests beginning in April
  - ▷ July: 24×7 support of this data path
- Filtering and Streaming
  - ▷ Level-3, Production from trigger table (using current Level-3 mechanism)
  - ▷ Offline format for detector/calibration data a priority (new Brandeis postdoc plus support from FNAL contract programmer)

Summary of Dates

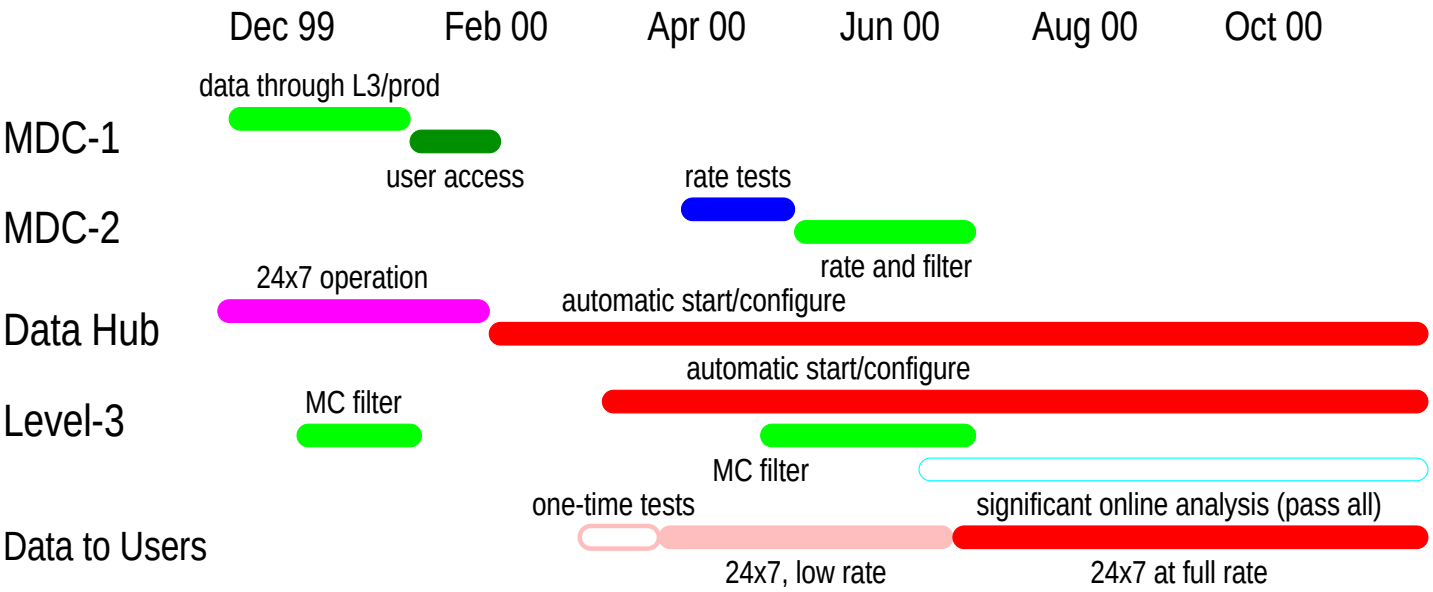

Conclusions

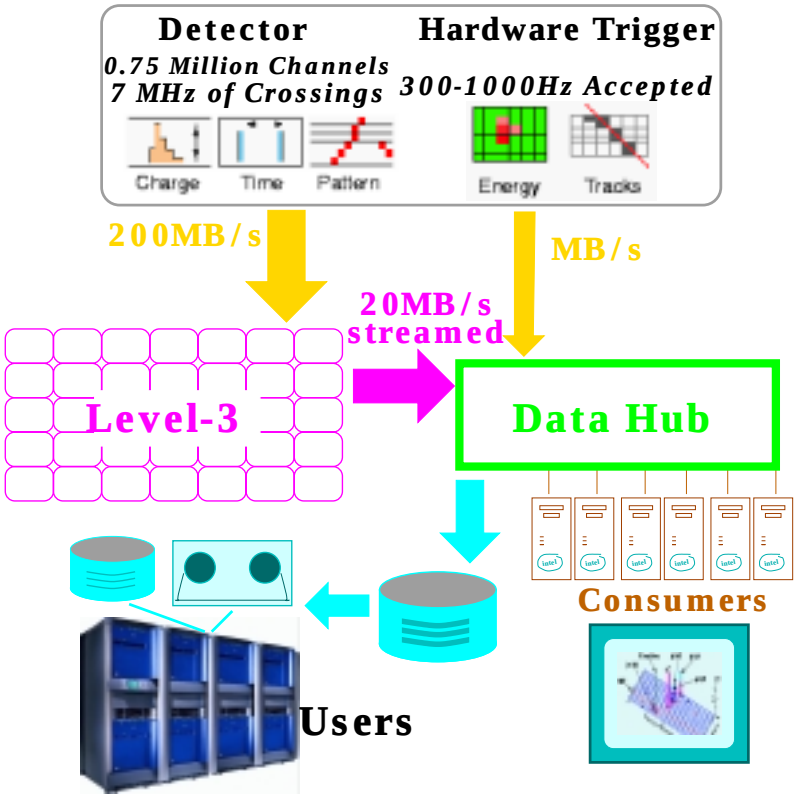

1. Bulk of hardware/software for high-level dataflow in place
2. Level-3 Filter Infrastructure functional
   - Reconstruction algorithms well underway
   - Heavy-flavor tagging algorithms and their ilk will require a learning curve with detectors and data
3. Data distributed to archive, consumers
4. Data handling tested, enter operations phase in March
5. Foci before Commissioning Run:
   - Adding last functionality
   - High bandwidth system tests
   - Convenient operations

Appendix: CDF-to-English Dictionary

Consumer Code that takes events from the *Data Hub* for online monitoring

COT Central Outer Tracker (open geometry wire chamber)

Data Hub The portion of the high-level DAQ that transfers events out of *Level-3* to *Consumers* and to the data logger. The truly horrific CDF Jargon for this machine is the "Consumer-Server/Data Logger"

EDM-2 The new CDF memory model, based on ROOT. Allows storage of object-oriented data. Replaces *TRYBOS*

Event Builder Portion of DAQ system that assembles events from different detectors into one place (*Level-3*). The "Software Event Builder" is a piece of low-rate code that emulates this functionality, bypassing *Level-3*.

FC Fibre-Channel. The network used for the disk stage of data logging.

FCC Feynman computing center; where the mass storage, production and central analysis system lives

FTE Full-time equivalent physicist. A ridiculously large time commitment of an individual physicist. Has been shown to alienate family and friends as well as result in population decrease.

ISL Intermediate (Outer) Silicon Layers

Level-3 The third Level of the CDF trigger system. Processor-based; run on a farm of low-cost PCs.

MDC Mock Data Challenge

Partition An unit of the DAQ system that can be independently controlled. Many independent partitions of the DAQ, through the *Data Hub*, can be run simultaneously for debugging.

Production The first data processing step after logging

ROOT The analysis and data management package that basis for the CDF II high-level physics analysis system and data format out of *Level-3*. Developed by Rene Brun's group at CERN.

SMP Symmetric Multi-Processor computer, e.g., SGI Origin 2000.

SVT Silicon Vertex Trigger. Level-2 trigger to identify impact parameters in Silicon, seeded from the *XFT* results.

SVX Inner Silicon Vertex Detector

Trigger Database The database that contains all trigger information required to define hardware and software triggers, data streams and production output streams.

TRYBOS The old (hardware) CDF event memory model, replaced by *EDM-2* in offline analysis.

XFT eXtremely Fast Tracker. Level-1 track-finding trigger in COT.