

Data Acquisition and Online Systems



Arnd Meyer
Fermilab

Collaboration Meeting Jan 24, 2002

All major components of the DAQ are operating
at design specifications or better, and are
capable of efficiently recording quality data

- Online Computing
- Run Control
- Front-End Systems
- Event Builder
- Level 3 Farm
- Consumer Server / Logger
- Consumer Monitors and Calibration Consumers



Online Computing and Security

Arnd Meyer
Jan 24, 2002

- The CDF online systems are one of the lab's critical systems. This meant additional requirements regarding computer security, and numerous network isolation, disaster recovery etc. tests
- As part of the CORC (**commissioning operational readiness clearance**) had to make a transition from an open development environment to a more secure environment for operations
 - Limiting access to the critical system while still maintaining the essential access. Access is controlled at the router through the access control list
 - Outlined plans for **data backup**. Critical disks are mirrored and are copied daily. Disks are backed up to tape weekly
 - Outlined a plan for the **failure recovery** of the various subsystems and spot tested the recovery procedure (including reconfiguration of backup hardware, reinstalling software, data recovery from backup, and booting from mirror disk). We have duplicate hardware for many of the critical systems.
 - Demonstrated ability to **take data while the BO network is disconnected** from the rest of the lab, including calibrations and monitoring.
 - Started an internal note documenting our recovery procedures
 - A number of items were identified and will be addressed for the next audit
- Expect the **next audit on the timescale of 6 to 12 months; unlike an ORC, the CORC will be periodically reviewed**



Online Computing

Arnd Meyer
Jan 24, 2002

- Additional Linux PCs on 3rd floor for Consumers and other server processes
- Upgrade to RedHat 7.1 is still in front of us (most of it), being evaluated
- Move web server away from important server (b0dau30)



Database

Arnd Meyer
Jan 24, 2002

- Overall working very well, no major crashes
- Progress highlights
 - Added about 240GB of additional mirrored disk space to online production, should be sufficient for Run 2a
 - New tape stacker for backup purposes is being installed
- To do (6 months)
 - Export to remote institutions is an ongoing project, biggest open issue
 - Implement web based histogramming of database quantities
- Performance
 - Continuous alarm monitoring works well in warning about both future and immediate problems (automatic emails and e-log entries)
 - Online data entry latency generally not noticeable (run control etc.), calibration latency much improved after data compression (QIE calibrations). L3Manager much improved, but still takes quite a long time
- Support and coverage
 - Have lost or soon losing many database group members: Dennis Box, Jim Kowalkowski, Kirsten Tollefson, Jack Cranshaw, Rick St. Denis
 - Excellent Oracle administrative coverage from highrise: Nelly Stanfield, Anil Kumar
 - Bill Badgett is new co-coordinator for database, concentrating on online; still looking for offline counterpart
 - All applications need ongoing support, all require expert training



Run Control

Arnd Meyer
Jan 24, 2002

- Running stably since some time now
- Continuous flow of requests for minor additions, most of them handled by a single person (Bill B.), with some help from Frank C., Mario, and myself
- Installed voice message server in October, can be used by any run control related process through smartsockets messages
- Retired old JPython scripts used in run control for online calibrations, moving more intelligence to the front-end crates
- Added additional monitoring for crate controllers, in response to power supply disaster
- Many additions to run database: intra-run rates, luminosity, losses, per-bunch info, etc., run summary with links to e-log, ...
- Much improved stability of ACNET interface to run control
- Significant improvements to error handling (Donatella T.)
 - Almost routinely using automatic recovery procedure for common failures (DONE timeouts etc.)
 - Will need to become more comprehensive and robust



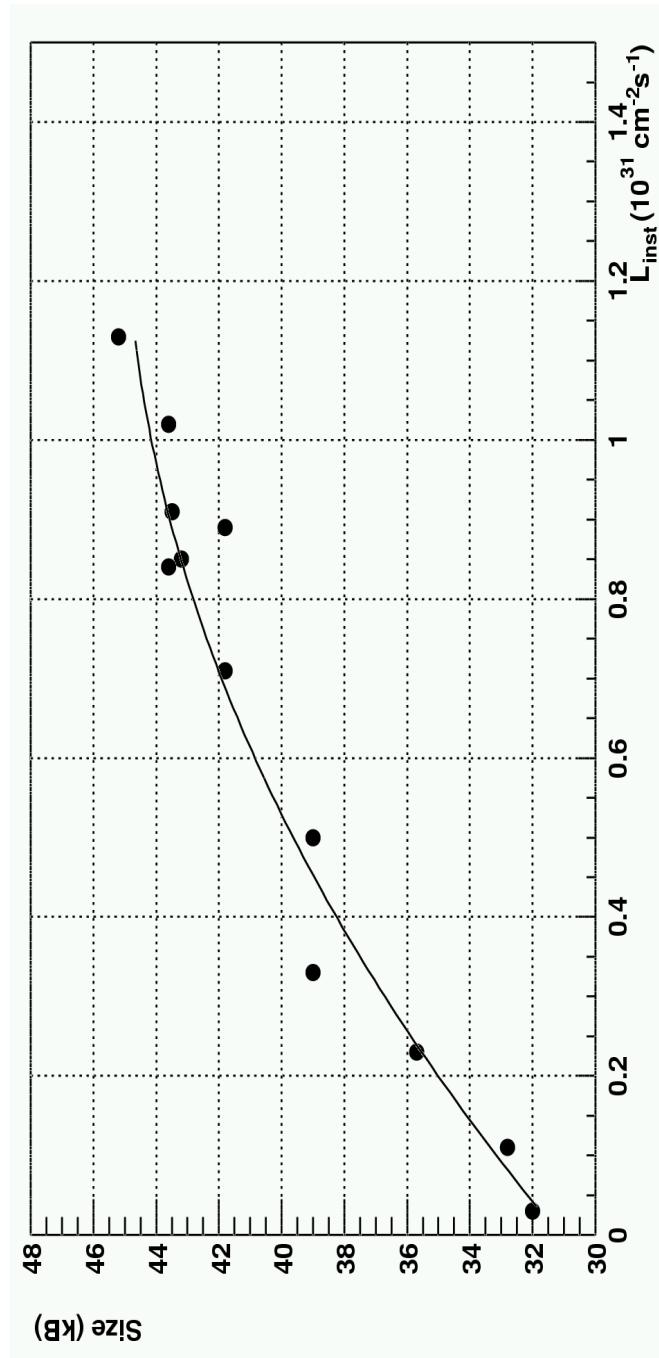
Front-End Systems

Arnd Meyer
Jan 24, 2002

- Front-End readout, initialization, and monitoring code is quite stable by now (thanks to / despite of ~40 people contributing!). Crates can stay up without resets for several months
- Continuous flow of requests for minor changes and bug fixes, mostly handled by one person (AM), with help from detector experts and DAQ personnel
- Recently an old friend, the "**crate lock-up problem**", came back. Symptoms have been analyzed, but no explanation. Workarounds exist, little impact on data taking
- Also occasional stability problems with SVT crates, under study
- Current maximum rate (L2A) around 500Hz, will go down with increasing luminosity
- We will start (again) investigating TDC spy mode readout
- Implementation of "diagnostic banks" to reduce data volume is waiting for Level 2
- Components added include COT FADCs, Miniplug, Roman Pots
- Implemented changes to front-end code to gracefully handle **CDF_ERROR**. Then it turned out that Silicon trips in HALTED state, so we had to disable the use of **CDF_ERROR** again. Also, more than 10 crates still have error lines disabled due to various problems

Front-End Systems: Outlook

- Limitations of some parts of the system are known since some time, and an upgrade path is detailed in the Run 2b TDR - e.g. event builder
- We took a careful look at the performance of front-end systems. Results are in CDF note 5824.
- Several systems will not be able to cope with the foreseen L2 accept rate of 1kHz, but software and minor hardware upgrades will suffice
- The one exception is the (COT) TDC system, which in its present form does not support rates beyond ~500Hz, and probably even less with increasing luminosity





Event Builder

Arnd Meyer
Jan 24, 2002

- Full configuration of EVB has been running for some time
- Several more VRBs will be added to the readout as the SVX group continues integration, no changes to EVB code required. Currently have >60 VRBs in readout
- Completed projects (last few months)
 - **Restructuring of the internal EVB/L3 network**, required for security and L3 farm size reasons
 - Debugging of EVB/ATM/L3 software for **large event sizes** (all known problems fixed)
 - Changes to VRB readout to overcome VRB firmware problems
 - **Current problems (none of them noticeably affecting physics datataking)**
 - There is ~**one hardware failure in several months** (ATM card in a SCPU)
 - Software: conflict with Silicon Readout Controller causes occasional BUSY timeout
 - Software: multi-partition running with multiple Halt-Recover-Run sequences causes EVB to crash after some time (obviously rare in "physics" mode)
- Near-term projects
 - Minor bug fixes
 - Work on EVB-SRC conflict, together with SVX group
 - Support for SVX monitoring project



Level 3

Arnd Meyer
Jan 24, 2002

- **Level 3 farm has grown to its final Run 2a size: 240 processor nodes, 16 subfarms**
- **Major completed projects (last few months)**
 - 120 PCs have been added in October/November
 - Fixes to dataflow code of Level 3; mostly related to reusing converters and outputs for subsequent runs
 - Significant improvements of reliability of the L3 monitoring
 - L3 transition times have not increased with doubling of the farm size
- **Current problems (none of them affecting physics datataking)**
 - There are **HD problems with the latest 120 PCs**, currently being evaluated
 - Occasional dataflow crashes
 - Automatic collection of filter log files for every run
 - A few additional minor problems are left
- **Near future projects**
 - Work on the above problems, more improvements to L3 monitoring
 - Study upgrade to RedHat 7.1



EVB / Level 3

Arnd Meyer
Jan 24, 2002

- In physics data-taking mode, running typically at 150-200Hz input to EVB/L3, 30-80Hz output of L3 with event sizes of about 200kB
 - No significant contribution to DAQ deadtime up to >300Hz input rate
- Reliability is good (1-2 pages/day), expected to improve further
- Monitoring and documentation are quite good
- Fully automatic interaction with L3 filter group (filter distribution and feedback on crashes)
- Joint projects on improvements will continue with SVX group
- Will have to revisit "reformatter errors" once Silicon is stable
 - Bookkeeping
 - Some event signatures might be more likely than others to be rejected
- Support
 - 2 grad students on EVB/L3 pager plus postdoc as backup, appears sufficient
 - Significant effort for system administration of ~286 nodes (total EVB+L3)
 - * MIT MOU specifies 1 FTE from FNAL/CDF for maintenance starting January 2002.
So far no person has been identified
 - * Minimal immediate needs: two half person days per week for L3 farm maintenance, to be trained by MIT group



Consumer Server / Logger

Arnd Meyer
Jan 24, 2002

- Is robustly handling data taking from multiple sources in multiple partitions
- Recently added 2 new disks at B0, increasing total online buffer space to 3.5TB
 - CSL is ready for increase from 4 to 8 raw data streams
 - Successfully tested last Friday
- Selection of events by trigger bit for Consumers is working since December
- Efforts to improve fast access to data
 - Investigating bandwidth upgrade for "look area" on fcdfsgi2 (dedicated Ethernet port, more disk space on fcdfsgi2)
- Monitoring, both for shift crew and CSL experts, works well
- To do for Run 2a:
 - Continue 24h pager support (Rochester, Tsukuba)
 - Code development to satisfy requests for improved/new functionality, e.g.
 - * Enhanced consumer interactions
 - * Specific requests to put a run onto tape (with Data Handling group)
 - * Investigate use of ENSTORE to write raw data to tape (CDF note 5822)



Calibration Consumers

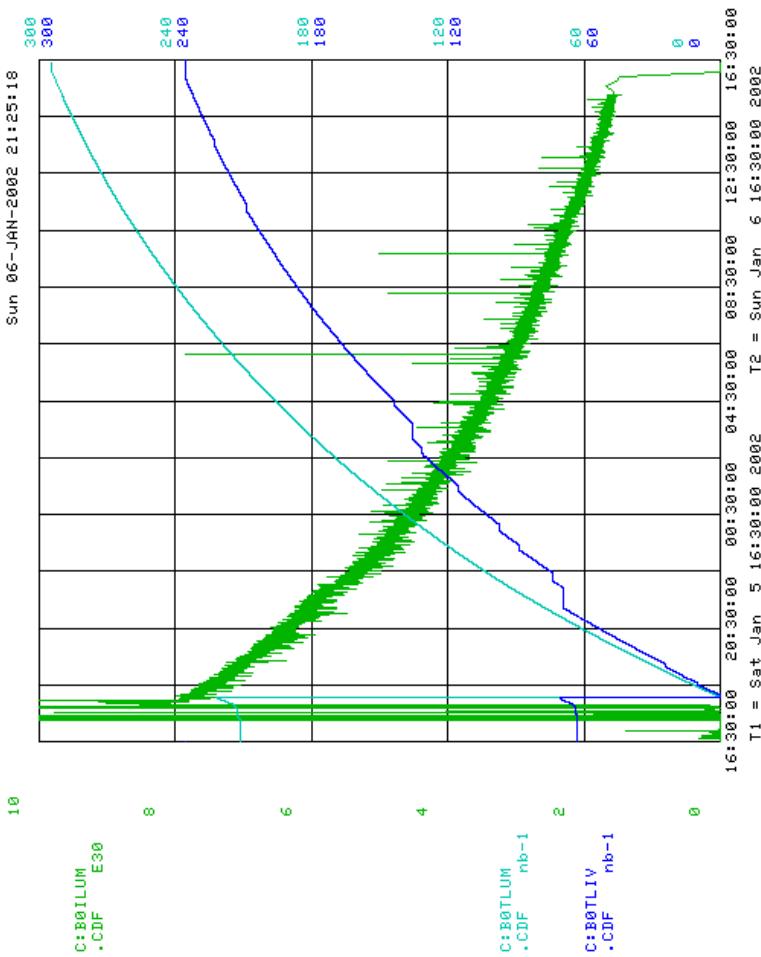
Arnd Meyer
Jan 24, 2002

- In general stable
- Should be upgraded to current consumer framework
 - E.g. for sending messages to run control, currently shift crews have to scan log files
 - Lack of manpower
- Also due to lack of manpower, maintenance of individual consumers has been turned over to detector groups
 - Need someone to make sure that people upgrade every now and then, so that we can remove old releases

Summary



Arnd Meyer
Jan 24, 2002



- 22 hour store with 80% efficiency
- Several shifts with >90% efficiency