DAQ Status

Arnd Meyer
Fermilab
Sept 19, 2002

Operating at design specs of 300 Hz L2 Accept rate, logging 20 MB/sec

- A single DAQ run of ~15 hrs with more than $0.5pb^{-1}$
- DAQ deadtime during this run 0.6%
- 97.8% of delivered luminosity recorded during this run

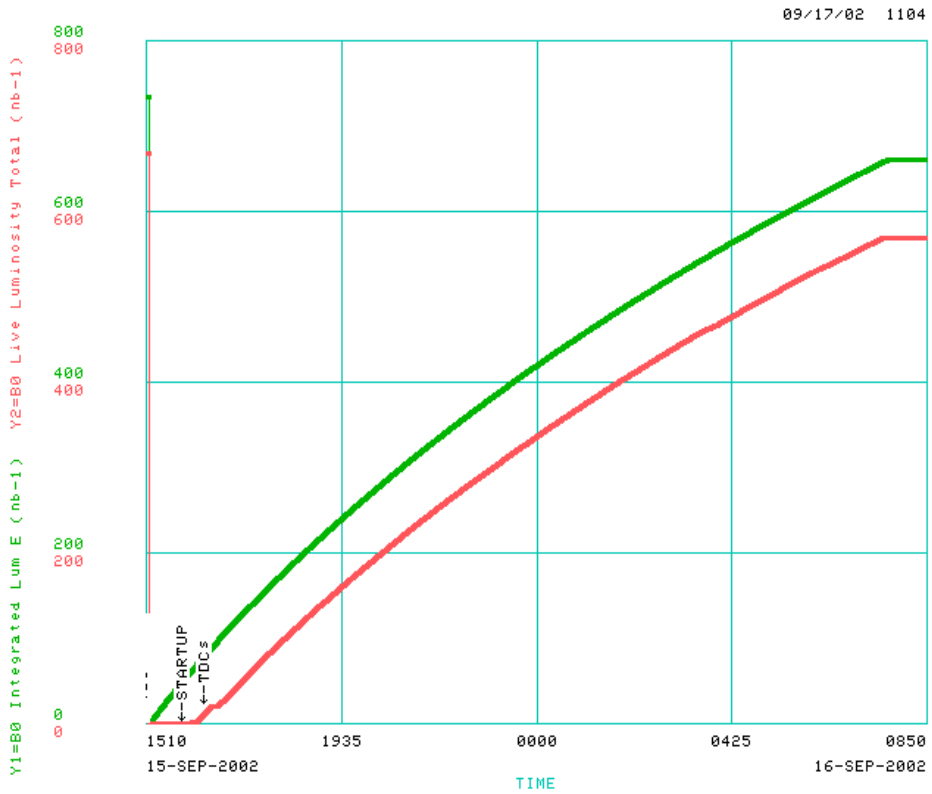

Run Control

Arnd Meyer
Fermilab
Sept 19, 2002

- New monitoring tools for shift crew, e.g.
  - ProcMon (important server processes), for early detection of potential problems, and faster recovery after power outages
  - TevMon, for quantitave evaluation of background situation, and early warnings
- Improved error recovery procedures (upgrades to state manager etc)
  - Much easier to recover e.g. from crate crashes
  - Reduce number of run-ends/run-starts
- Error handler undergoing continuous changes to better assist the shift crew
  - Had introduced a couple of annoying bugs (like opening hundreds of windows on the same error...)
- Under development are tools for DAQ downtime accounting
  - Log causes for downtimes, failures, HRR's etc.
  - Report in terms of lost integrated luminosity
  - ACNET downtime logger not very useful for this: no correlation to DAQ events
- Dynamic prescaling ready to go, final adjustments underway
  - Fully exploit available bandwidth

Front-End Readout Progress

Arnd Meyer
Fermilab
Sept 19, 2002

- TDC *"Spy"* mode routinely used since July (after fixing TRACER firmware in all COT crates); last missing major design feature of TDC's
  - As a side effect, don't truncate COT data in the front-end crates anymore
- Made calorimeter (ADMEM) bank formatting significantly (x2) faster
  - Major slow-down is now central SMXR (and to a lesser extent Plug)
- Successfully tested new crate CPU (Motorola MVME 2400 series), for selected processor upgrades. Ordered 10, to be installed in crates which benefit most.
- Trigger Supervisor Readout Lists ready to go
  - "diagnostic" banks read out only on special L2 triggers → faster readout, reduced data volume)
  - Required concerted effort of TSI, L2, DAQ, TDWG groups
  - XTRP, XFT Finder are among the slowest crates presently. L2 decision will require some more thought (will need this for TL2D due to slow VME transfers of Alpha).
- Overall code-optimization, improvements, additions (e.g. TOF Trigger), etc.
- Protection against high trigger rates (or other dangers) for Silicon

Readout Status

Arnd Meyer
Fermilab
Sept 19, 2002

Time required until a given L2 Buffer is freed after L2A

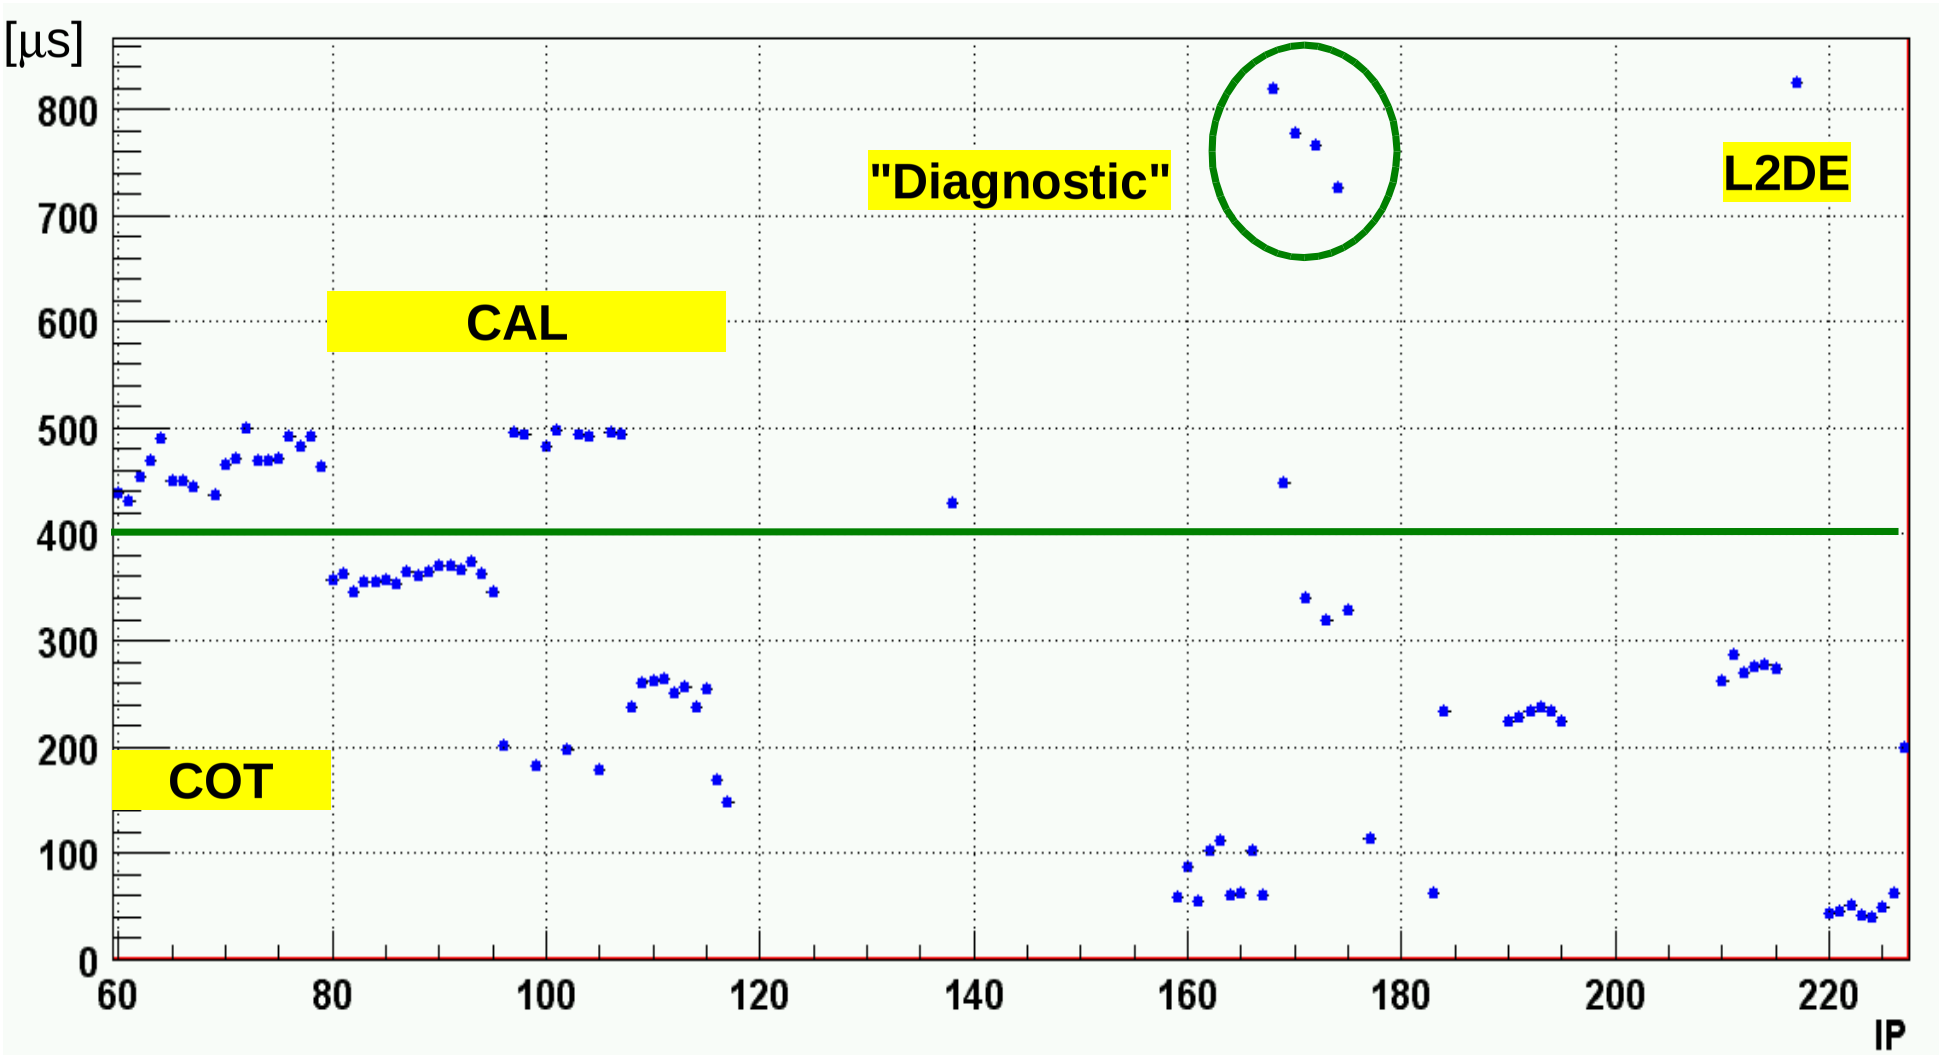

Run 151476, L~2.5·10^{31}cm^{-2}s^{-1}

Readout Status

Arnd Meyer
Fermilab
Sept 19, 2002

Time for complete readout cycle

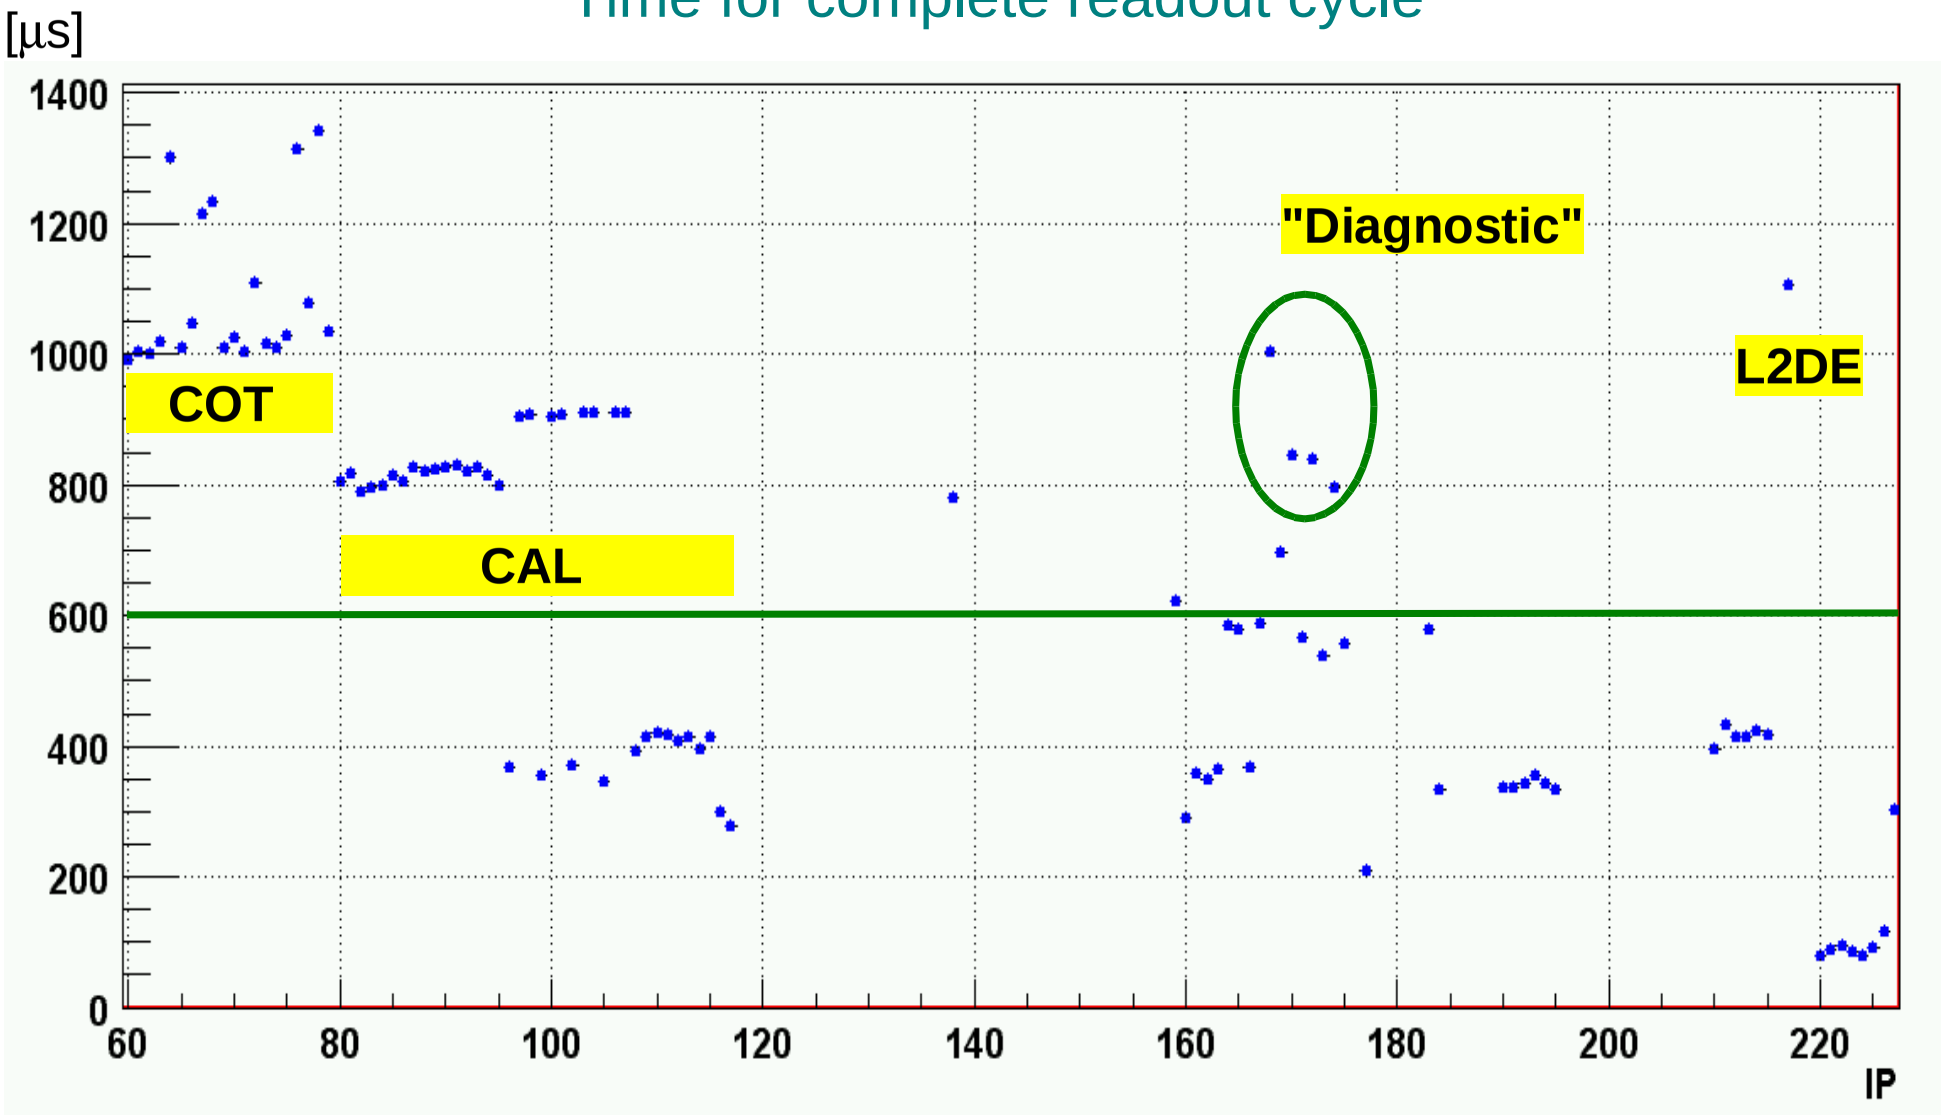

Outstanding Front-End Issues

Arnd Meyer
Fermilab
Sept 19, 2002

- Crate crashes: rare, most crates can run for months without a problem. However, the two "fcal" crates (Miniplug, Roman Pots) are notorious.
- Although front-end problems are not a DAQ problem, their impact on the system is ⇒ improve robustness, diagnostics, instructions
- TDC's "hanging": will likely have to live with this at some level
- One COT crate not yet read out in *"Spy"* mode, because of increased rate of data corruption
  - Attempted to swap TRACER, use spare fiber, switch BUSY and DATA fibers, swap VRB transition board, swap VRB → problem still observed
- Large number (>10 / shift) of front-end errors that require a *Halt-Recover-Run* sequence to recover.
  - Typically < 2% of luminosity lost due to these, but annoying
  - DONE timeouts because of "SMXR not responding", bunch counter errors, ...
  - Level2 decision timeouts
- Clock errors cause massive amount of errors, usually requires start of a new run. Basic clock monitoring is in place.
- Minor nuisances: BC errors from XFT Finders, VME errors from TOF trigger, etc.

“VRB Problems”

(CP: “This is not a problem of our group”)

Arnd Meyer
Fermilab
Sept 19, 2002

- DAQ VRB's collect data from 10 FE crates each, installed in 1st floor EVB crates
- Somewhat neglected in the past; e.g. couldn't find spare VRB transition module, first one found didn't work
- VRB bus errors (“Can't reset VRB” error)
  - On average once per day, a VRB goes into a state where it is unresponsive to VME.
  - Until ~10 days ago, required powercycling the crate and subsequent reinitialization→ lengthy, ~15min or more depending on shift crew experience
  - Developed new and faster procedure, just resetting the affected VRB→ saves a lot of time. Simultaneously, working on a real solution (new VRB firmware, ...)
- Corrupted VRB headers: causes frequent *Halt-Recover-Run's*, occasionally EVB dies. Again, try new firmware, replace VRB's, ...
- Low level “data corruption errors” from various crates (swapped bytes etc.)
  - In most cases rare enough not to be a major worry (once/day)
  - Can have many sources → hard to diagnose
- “TRACER word event ID problem” (up to once/30min): not understood. Requires HRR, with a potential of EVB crash. Not necessarily a VRB problem.

Event Builder

Arnd Meyer
Fermilab
Sept 19, 2002

- Primary concern are incidents where the event builder is "stuck" → end run + *"clean up event builder"* + restart takes 10min (can be much more if problems occur)
- Under normal running conditions, happens of the order of once per day
- "Events with missing banks": under certain fatal error conditions, only parts of some events make it to L3. Will be fixed with new version of EVB software (reject complete event).
- L3 subfarms "freezing" in the middle of the run: problem appears at the interface between L3 converter nodes and EVB. Hope is that better bookkeeping of event buffers in SM/SCPU's will solve this. Hasn't happened recently (last week).
- EVB crash during *Halt-Recover-Run*: happens every now and then.
- A new version of the EVB software is under testing now, and is supposed to address all of the above

Level 3

Arnd Meyer
Fermilab
Sept 19, 2002

- Level3 node failures (*not unexpected* with ~250 PCs)
  - Usually detected and gracefully handled by control code: bad nodes are subsequently ignored, and a warning is sent to the aces
  - Unfortunately the monitoring that warns the ace sometimes died. Patched now.
  - Sometimes L3 gateway crashes, requiring expert to reboot it (usually more than 1 hr data lost)
- Reformatter code update: recovering now most events with minimal Silicon data corruption
- Long transition times (Partition/Coldstart)
  - Starting a run takes about 5min
  - Level 3 is the single largest contributor
- Severe filter errors
  - Many caused by high COT occupancy
  - Will add rate of errors to ace tools
  - One probable cause for larger error rate: COT data are no longer truncated

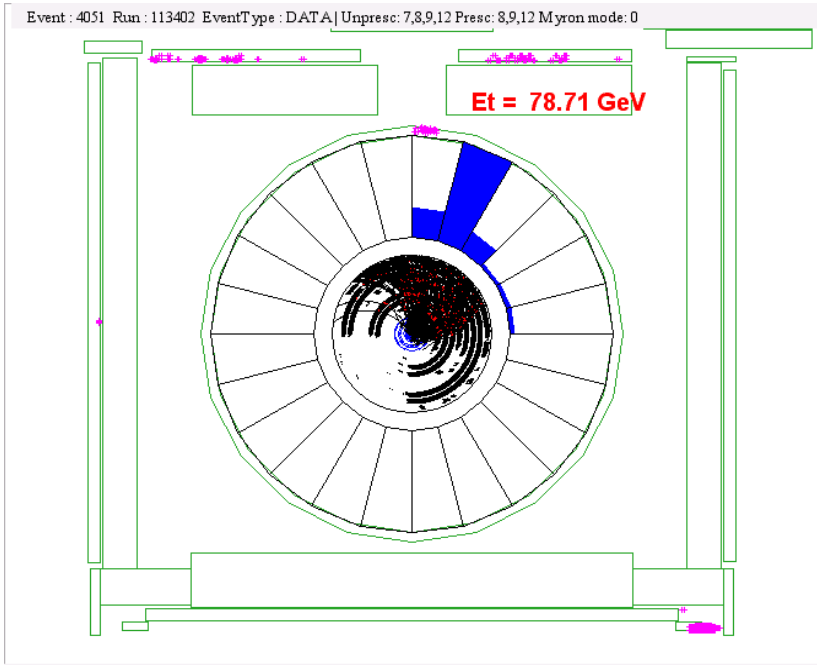

April 2001

Consumer Server / Data Logger

Arnd Meyer
Fermilab
Sept 19, 2002

- In general rarely noticed on shift → good!
- Just two days ago, lost almost 2hrs due to CSL failure
  - Failure is understood (memory problem), and a patch is underway
  - Also monitoring was not working. Will add error messages, improve monitoring
- Wishlist
  - Understand rare failures
  - Improve (shift crew) monitoring to faster recover from failures
  - Improve robustness of “calibration CSL”: sometimes gets into a bad state when confronted with high rates

Summary

- DAQ systems are operating efficiently, routinely collecting 80-90% of the delivered luminosity
- Most causes for “long” downtimes have been eliminated
- Addressing the long list of minor problems will incrementally improve data taking
- Will always have some level of hardware failures, procedures to deal with those quickly are important