Fermi Gamma-ray Space Telescope

Design and implementation of the Fermi LAT level 1 pipeline

Warren Focke

SLAC SCA

Orbital gamma-ray telescope
- Launched June, 2008

We want to turn the data around quickly
- Because something might be happening right now

But it's in orbit
- Which means data delivery to us is bursty and, well...
  - Sometimes the data arrives out of order
  - Sometimes they send it twice (or more)
  - Sometimes it arrives late
  - Sometimes it never comes (this is rare)

Data Flow

TDRSS
- NASA's near-Earth space network

White Sands
- TDRSS base station

NASA Goddard
- Fermi Mission Operations Center

SLAC
- Fermi LAT Instrument Science Operations Center
  - Up to here it's been data straight out of the spacecraft recorder.
  - Now we turn it into photons and send them to...

NASA Goddard
- Fermi Science Support Center

Processing

- Takes ~125 hours of CPU to process 1 hour of data
  - Plus waiting for disks, network, ...
  - So ~150-175 core-hours
- Data is delivered ~ every 3 hours (3 hours worth)
  - And the project wants that processed in an hour
    - Which we almost manage
  - So we need ~500 cores
  - OCIO gives us 800
- Where does the time go?
  - Mostly reconstruction
  - Plus lots of instrument health monitoring
  - 2000-3000 jobs every 3 hours
    - About half of them are recon
- Multiple levels of parallelism
  - We try to make most jobs take 5-30 minutes
  - Change block sizes, shuffle jobs between levels to achieve this
- This is mostly Python wrappers around other people's code
  - Including a library of things useful to most pipeline users that's shared with MC, reprocessing, and Automated Science Processing
    - GPLtools – ask me or Tom Glanzman

Dependency Flow

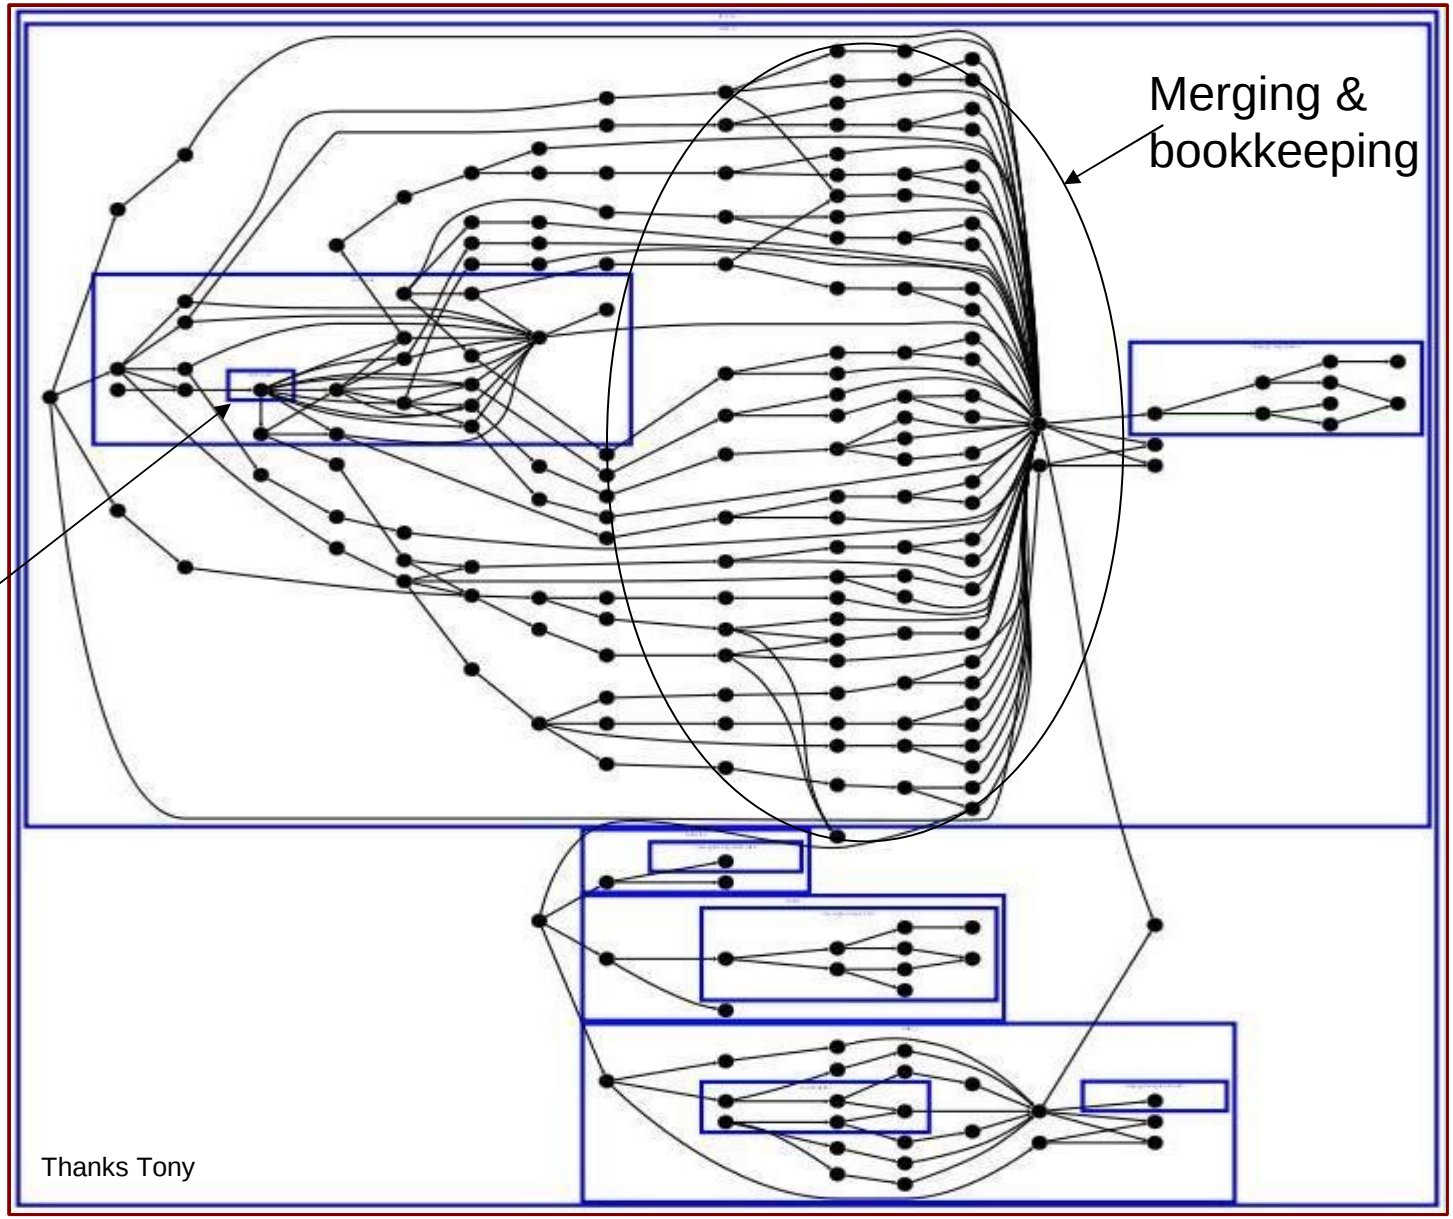

Performance vs Time

Misplaced 4 hours – found it 2 years later

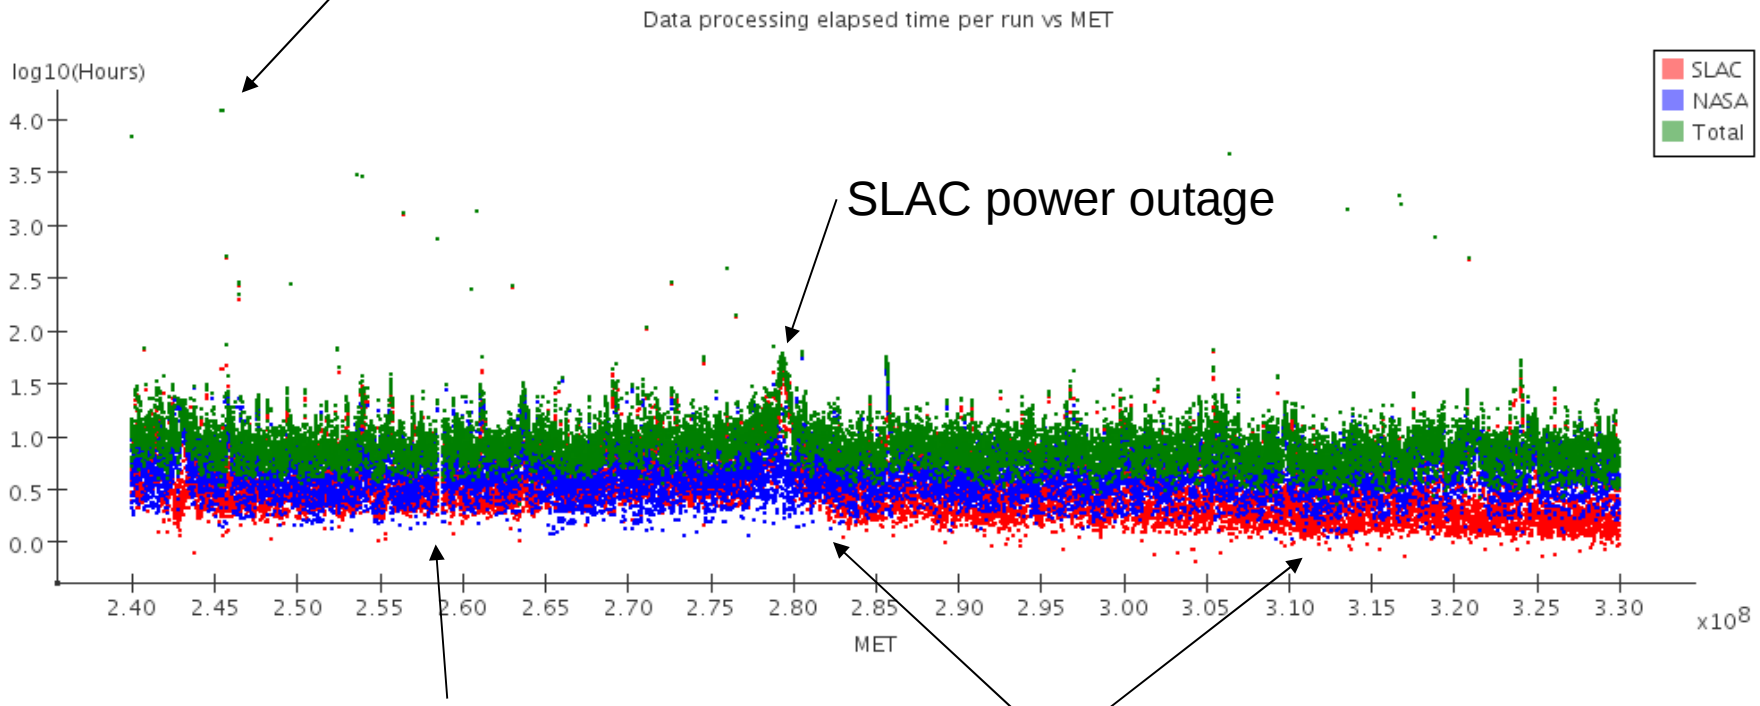

SLAC power outage

Load shed - SEU?

I/O upgrades