CSA07 Issues

Ian Fisk

November 28, 2007

CSA07 Workflows

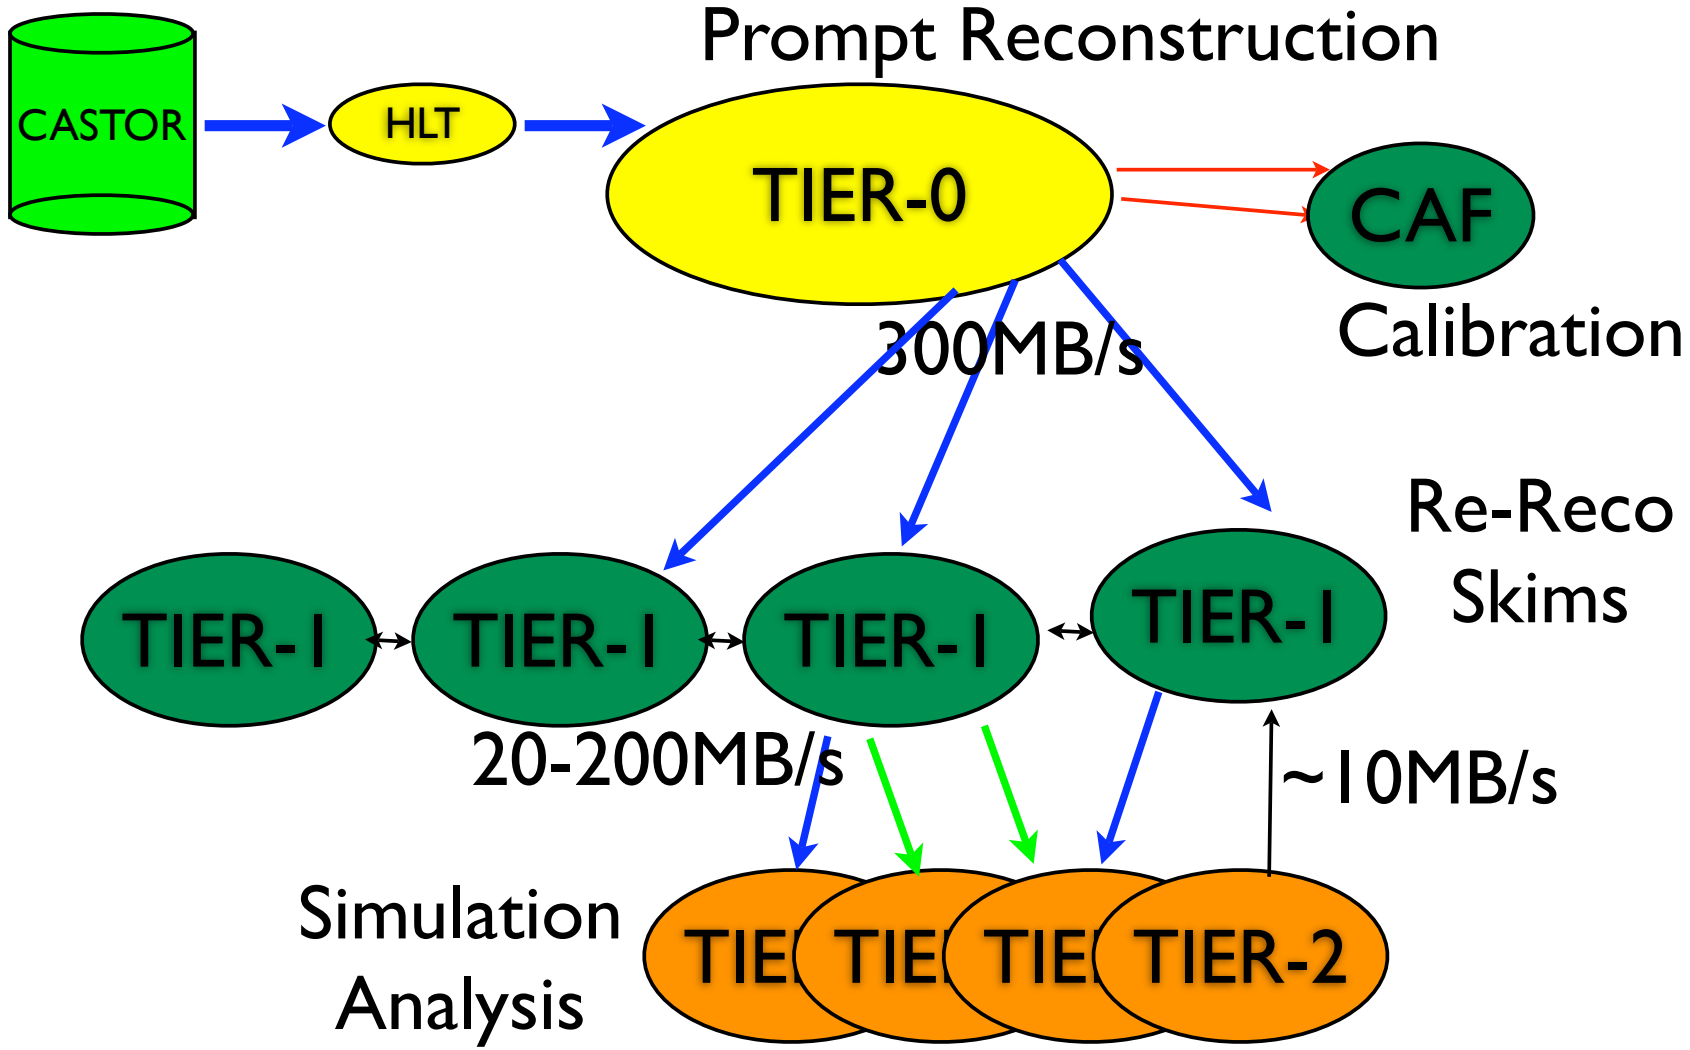

Tier-0 to Tier-1 Transfers

In the CMS model the Tier-1 centers receive raw data

- A natural extension of the online of the experiment
- Data is not considered "safe" until the second copy is on tape
- CERN is the "cold" copy served data comes from Tier-1s

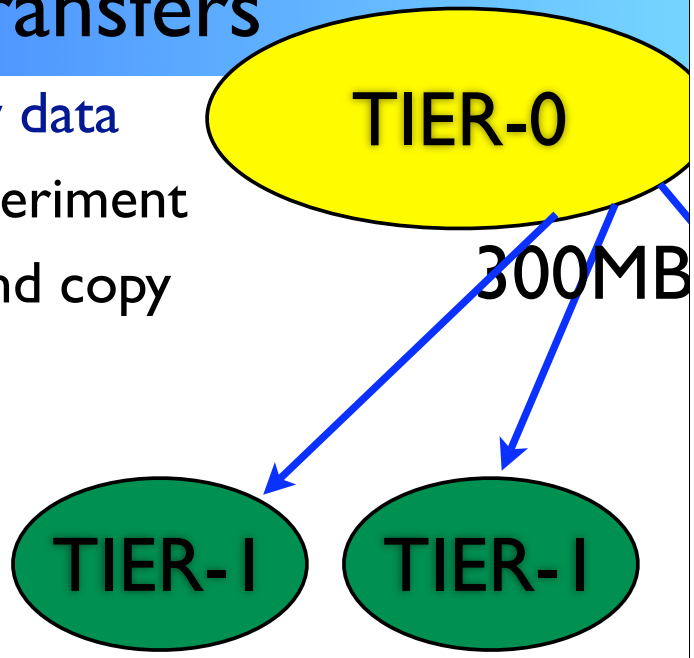

One item we didn't exercise well in CSA07 was to check the latency to get the second copy of the data exported and received at the remote Tier-1

- Something to work on between now and February for the functionality tests in CCRC
- The automated handling of the injection and subscription process needs to be improved.

Transfers From CERN to IN2P3

We have successfully met the transfer metrics,

- Rate from CERN is not as routinely high as we expected
  - Data was not available for transfer at all times

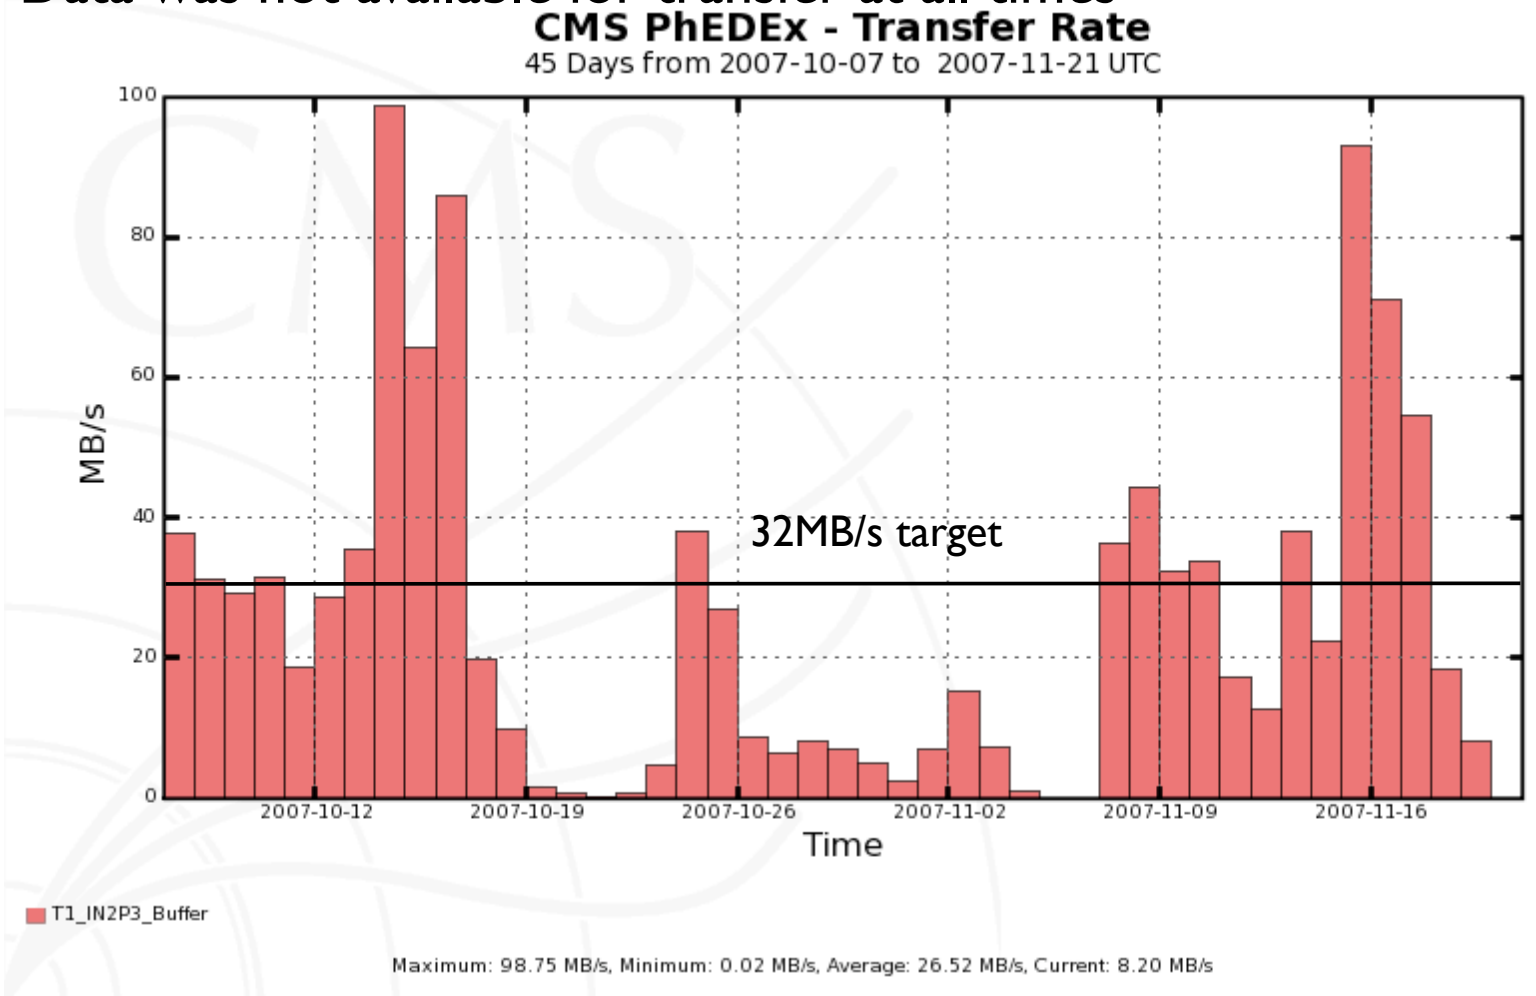

Transfer Quality

Transfer quality is not as green as as we might like

- There are two end points for each transfer, and in some cases we have stressed both of them
  - May point to the need partition

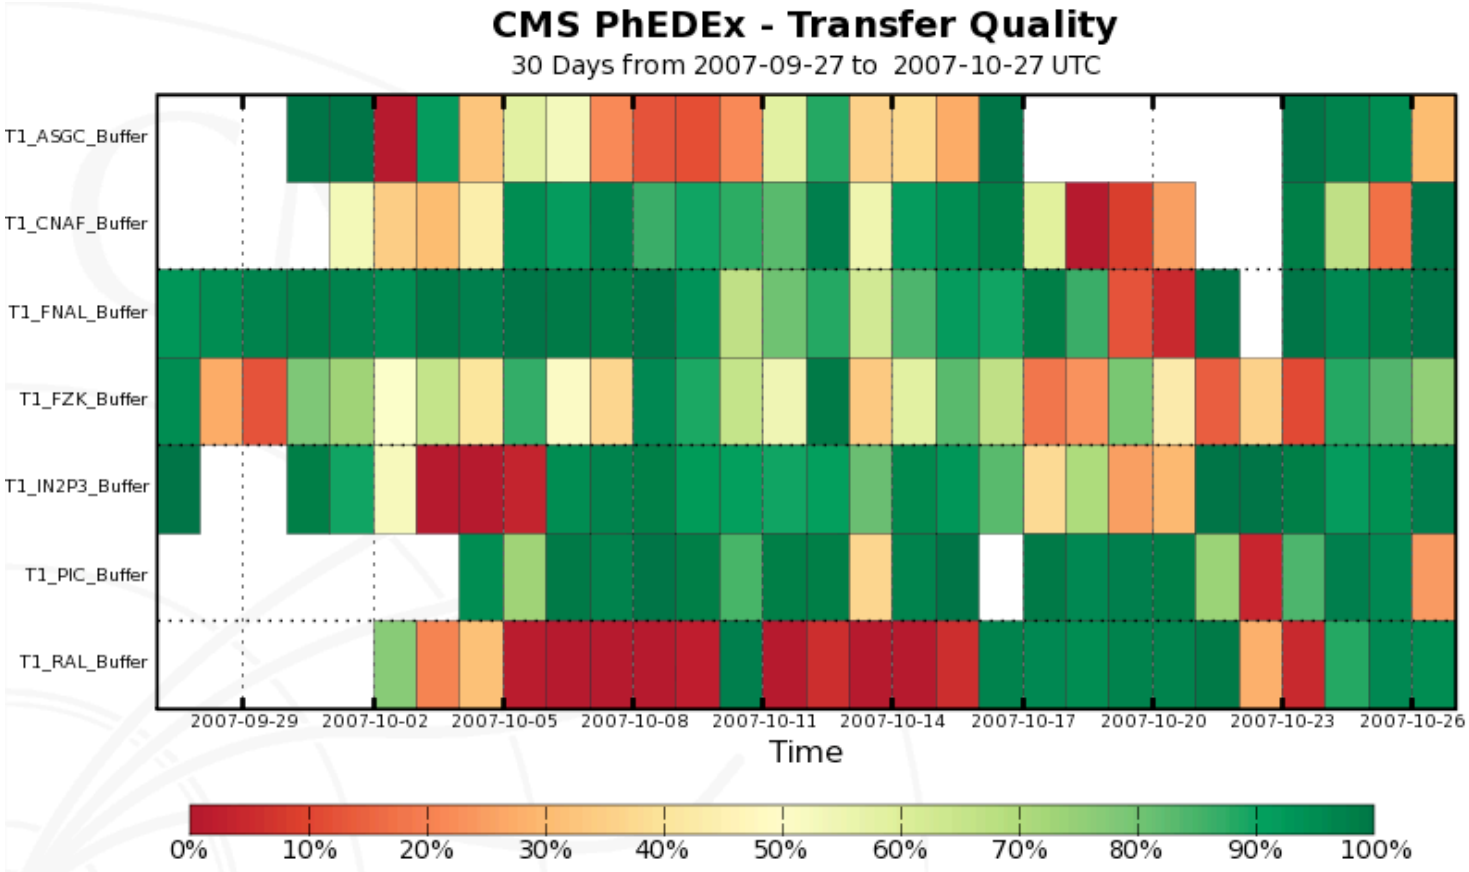

Transfer Quality all Sites to IN2P3 Production

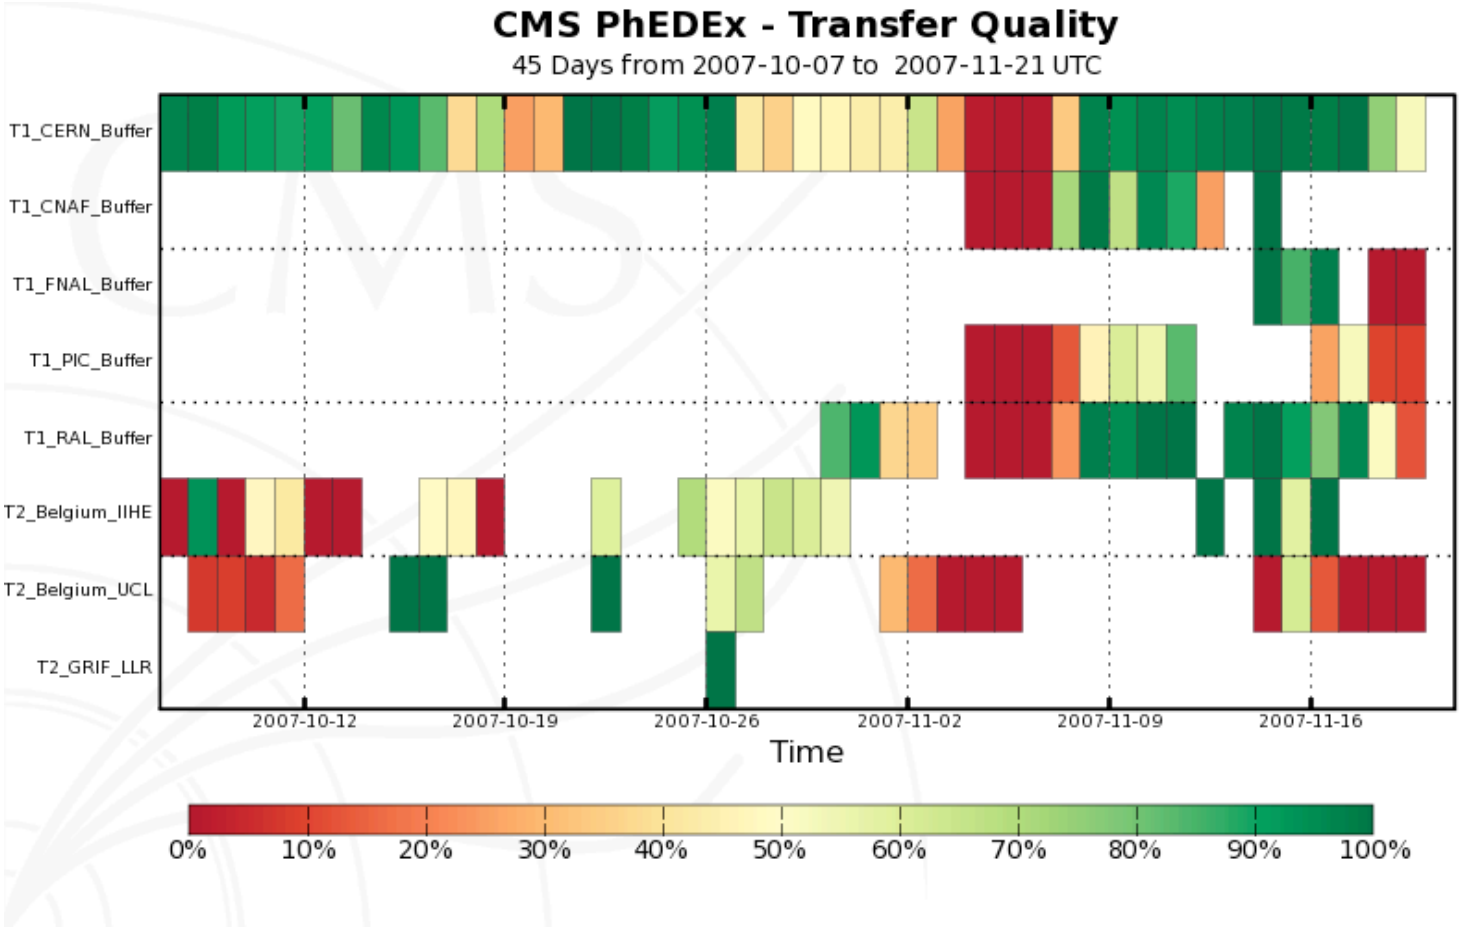

IN2P3 Production Transfers

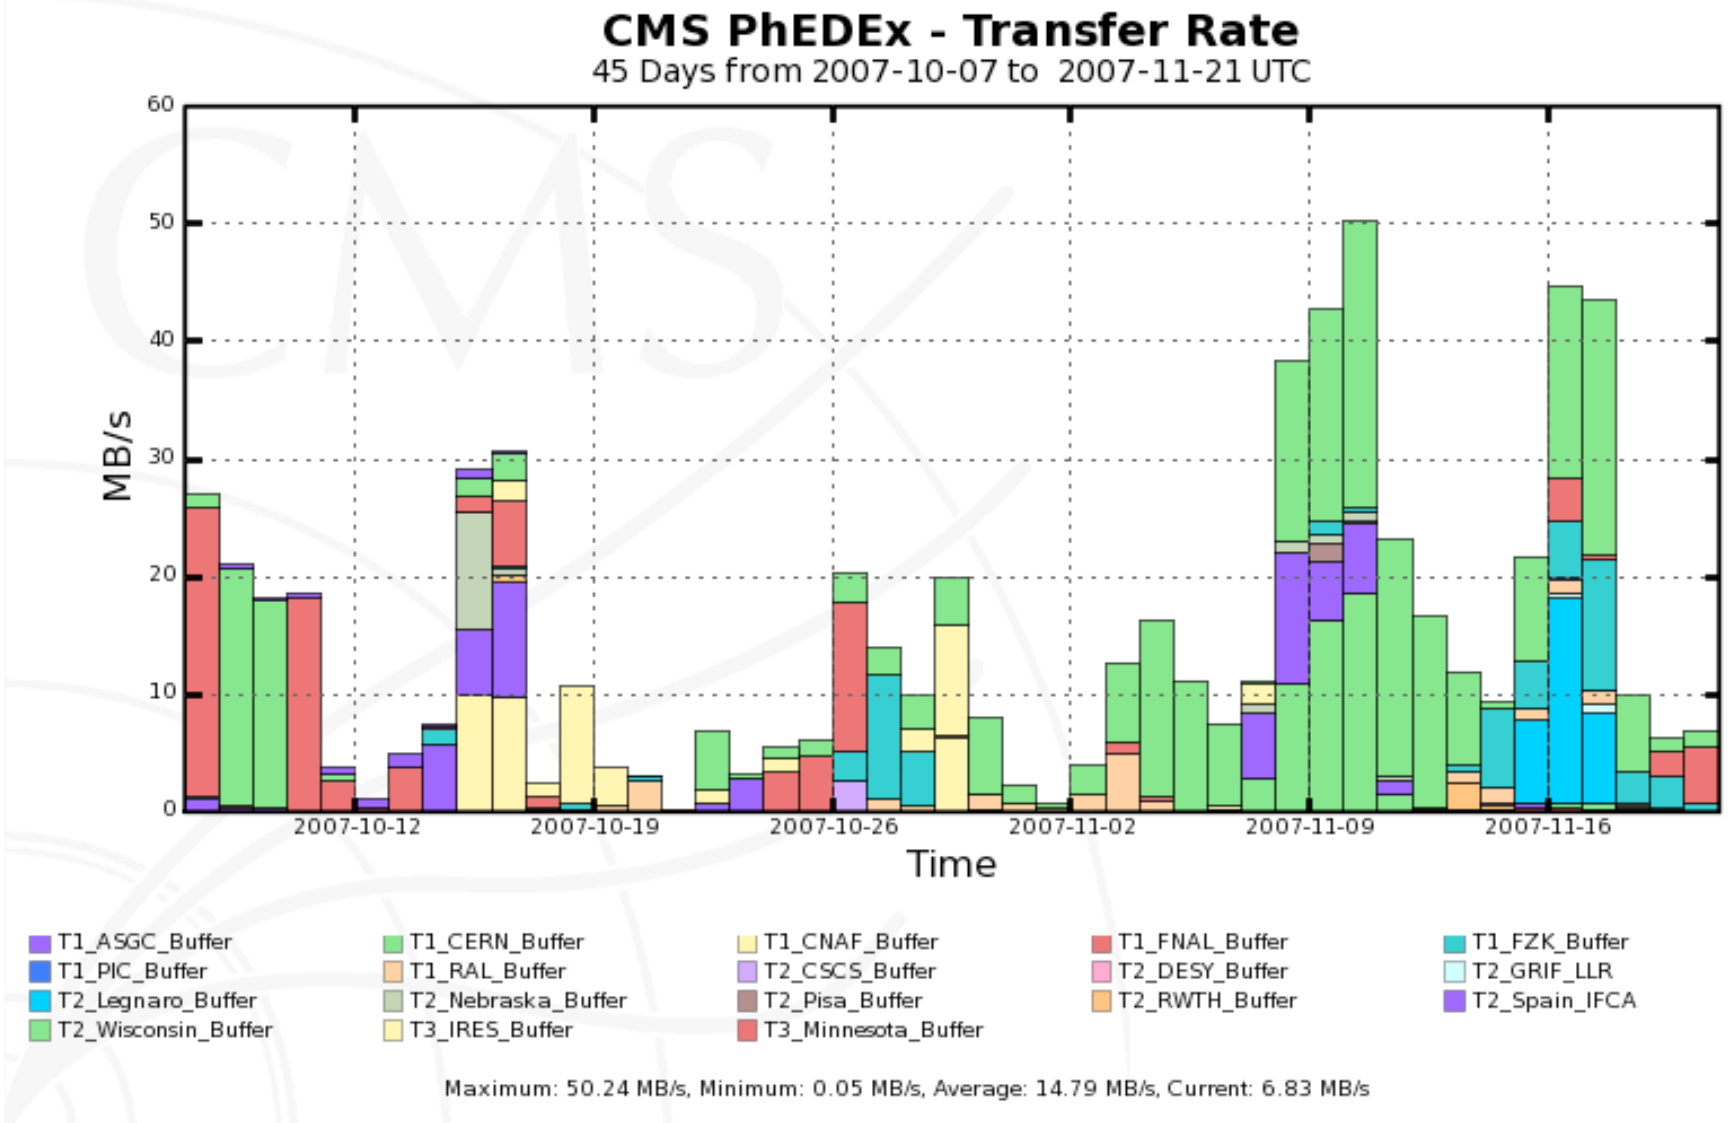

Transfers From IN2P3 Debug

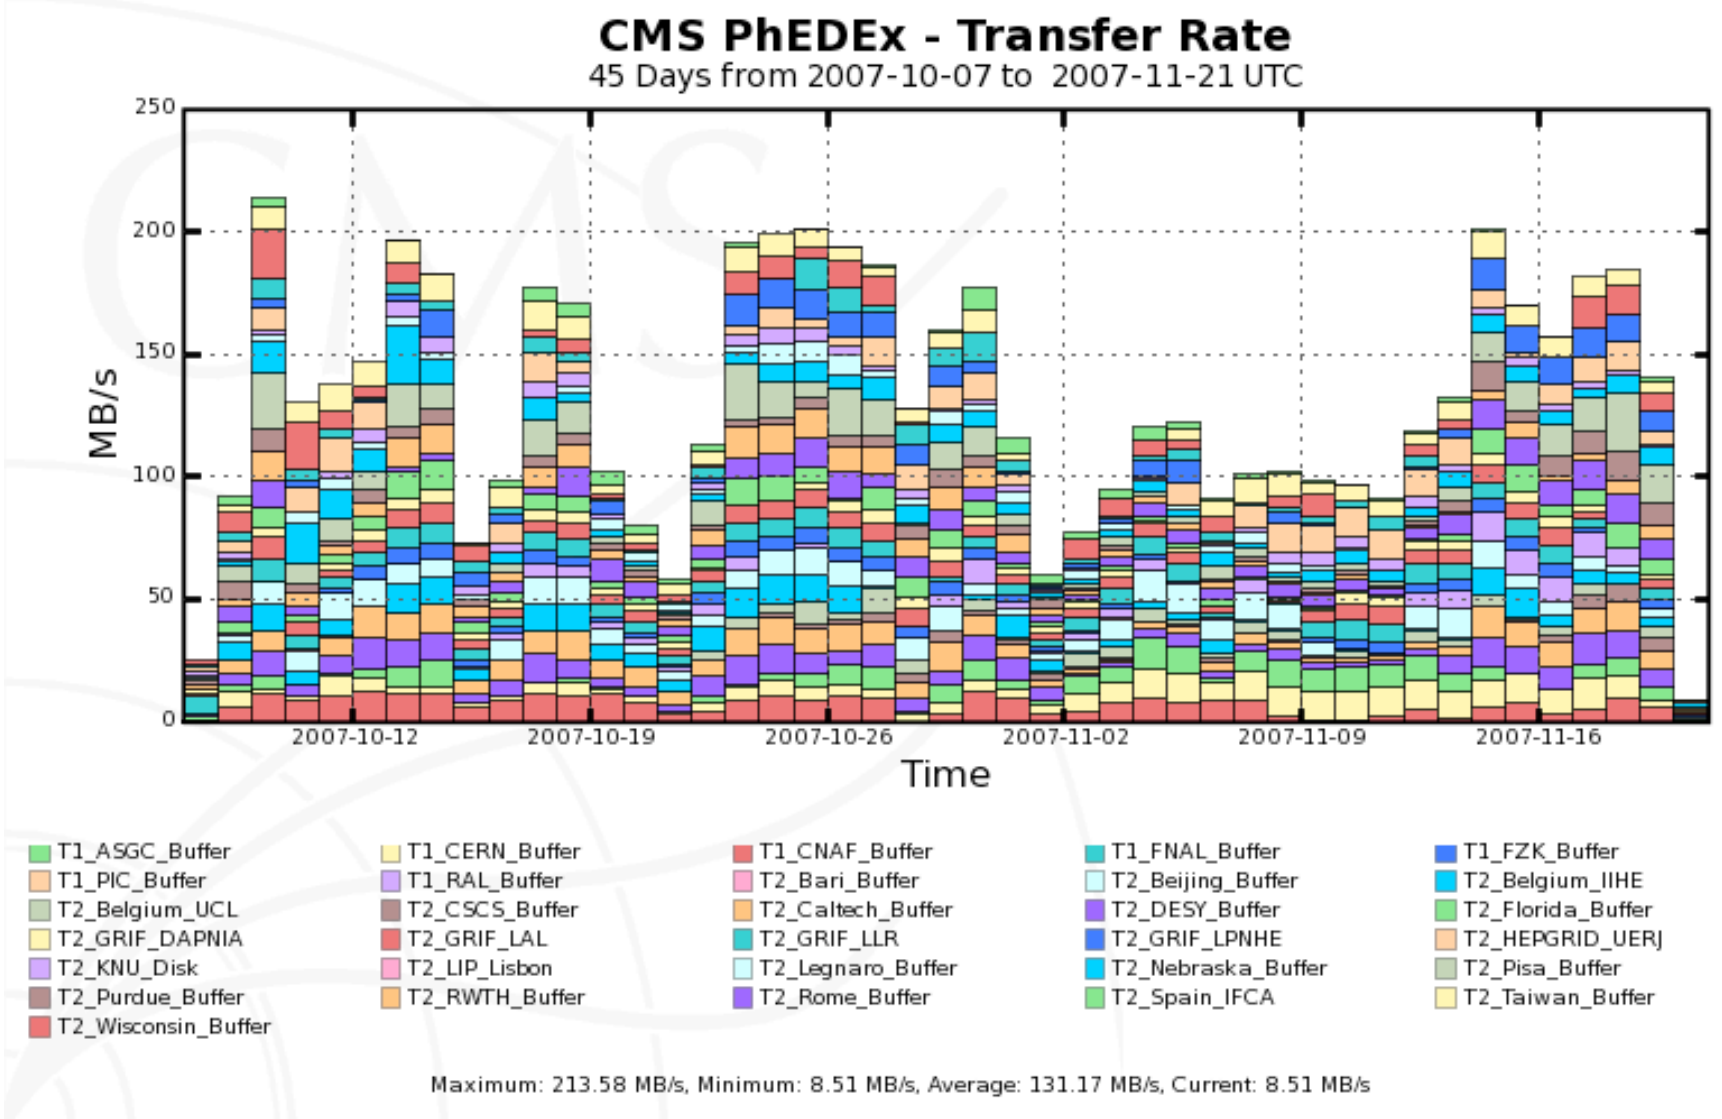

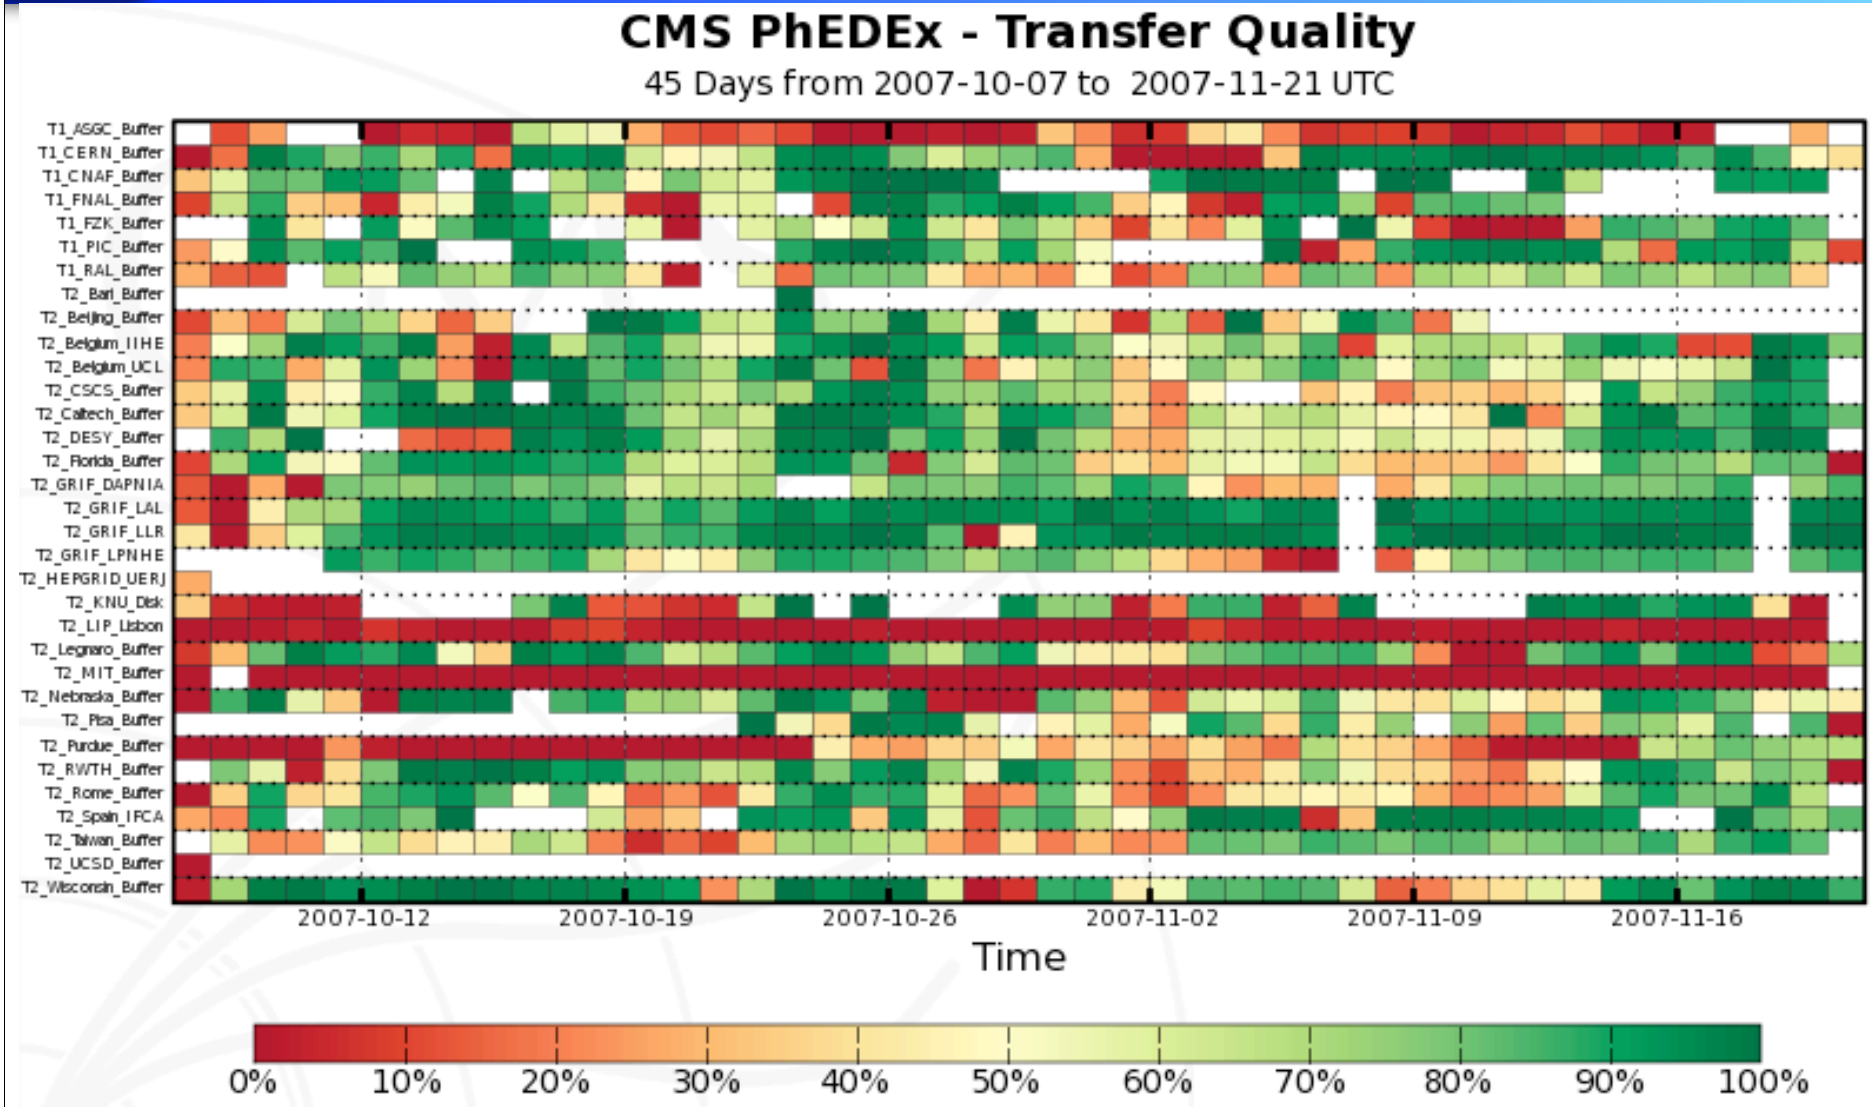
CMS PhEDEx - Transfer Quality
45 Days from 2007-10-07 to 2007-11-21 UTC
T1_ASGC_Buffer
T1_CERN_Buffer
T1_CNAF_Buffer
T1_FNAL_Buffer
T1_FZK_Buffer
T1_PIC_Buffer
T1_RAL_Buffer
T2_Bari_Buffer
T2_Beijing_Buffer
T2_Belgium_IIHE
T2_Belgium_UCL
T2_CSCS_Buffer
T2_Caltech_Buffer
T2_DESY_Buffer
T2_Florida_Buffer
T2_GRIF_DAPNIA
T2_GRIF_LAL
T2_GRIF_LLR
T2_GRIF_LPNHE
T2_HEPGRID_UERJ
T2_KNU_Disk
T2_LIP_Lisbon
T2_Legnaro_Buffer
T2_MIT_Buffer
T2_Nebraska_Buffer
T2_Pisa_Buffer
T2_Purdue_Buffer
T2_RWTH_Buffer
T2_Rome_Buffer
T2_Spain_IFCA
T2_Taiwan_Buffer
T2_UCSD_Buffer
T2_Wisconsin_Buffer
2007-10-12
2007-10-19
2007-10-26
2007-11-02
2007-11-09
2007-11-16
Time
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%

Organized Processing

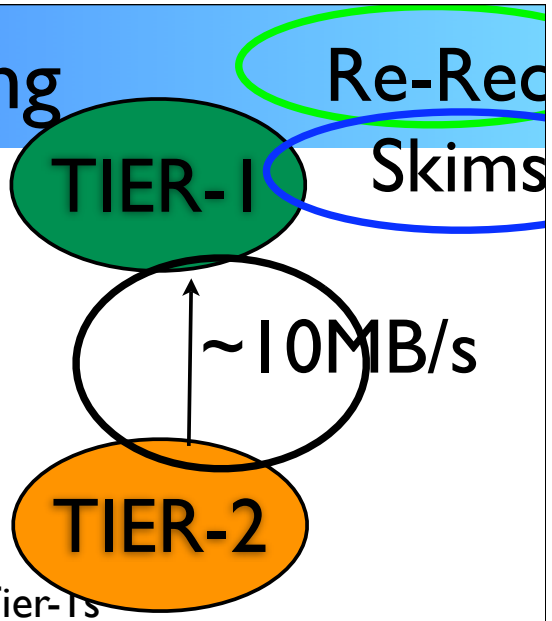

Tier-1 centers in the CMS model perform organized processing of the data entrusted to them

- Re-reconstruction
  - In CSA07 this included both incoming "data"
  - Simulated "signal" samples from Tier-2s
    - Signal samples were reconstructed with high priority at Tier-1s
- Skimming
  - Skimming came late
    - A lot of skims are available.

Process Submissions

IN2P3 jobs scale as one would expect

- Roughly 2200 jobs per day

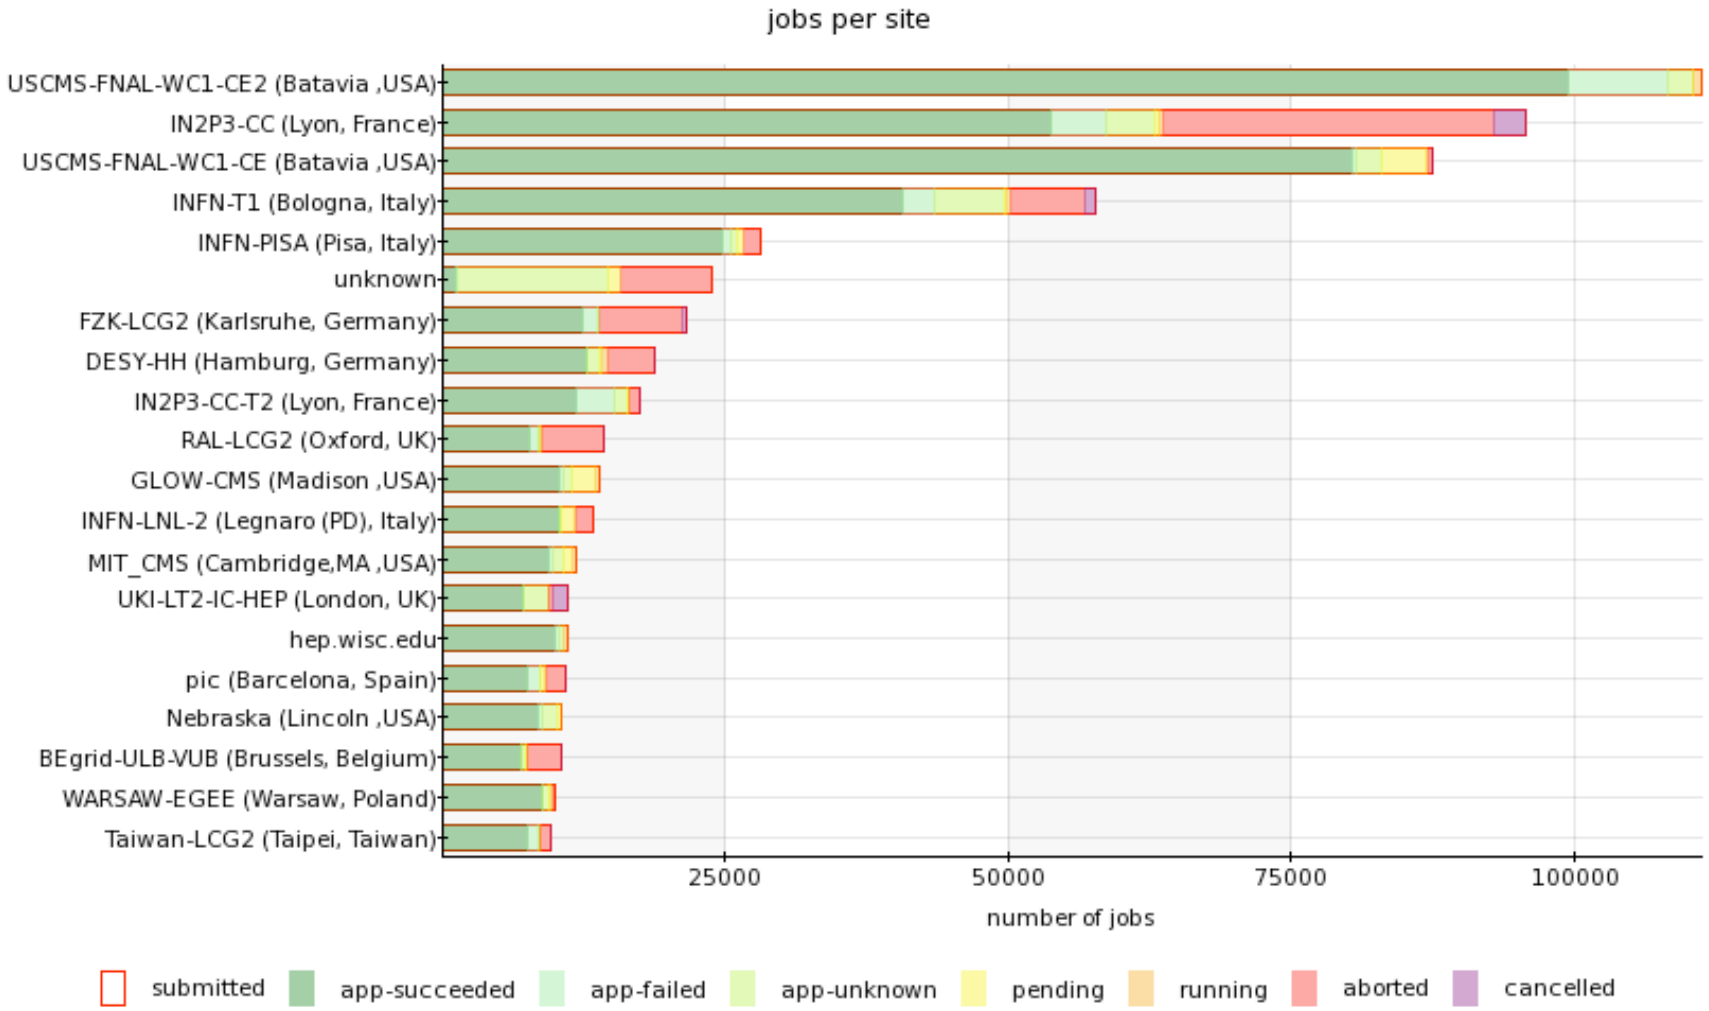

Processing Concerns

Merging issues

- A lot of the was CMS software problems, but we should understand if the problems hit IN2P3 harder than other places

We have run skimming and signal reconstruction

- Next step is to run reprocessing with various calibrations.
  - Can exercise mass storage

Other upcoming work

One other item for PADA is to increase the analysis submission rates

- Need to work with the Tier-2 sites in the region to prepare for analysis submissions
- Energizing the local communities to use the sites for analysis

Need to work on the perceived error rates in CRAB to support distributed analysis

- Work to validate the tools on the local sites and to support for users

Analysis Submissions

There were some untracked submissions over the month of the challenge

- Could be monitoring of site problems

IN2P3 and Beligum figured prominently in the analysis plot

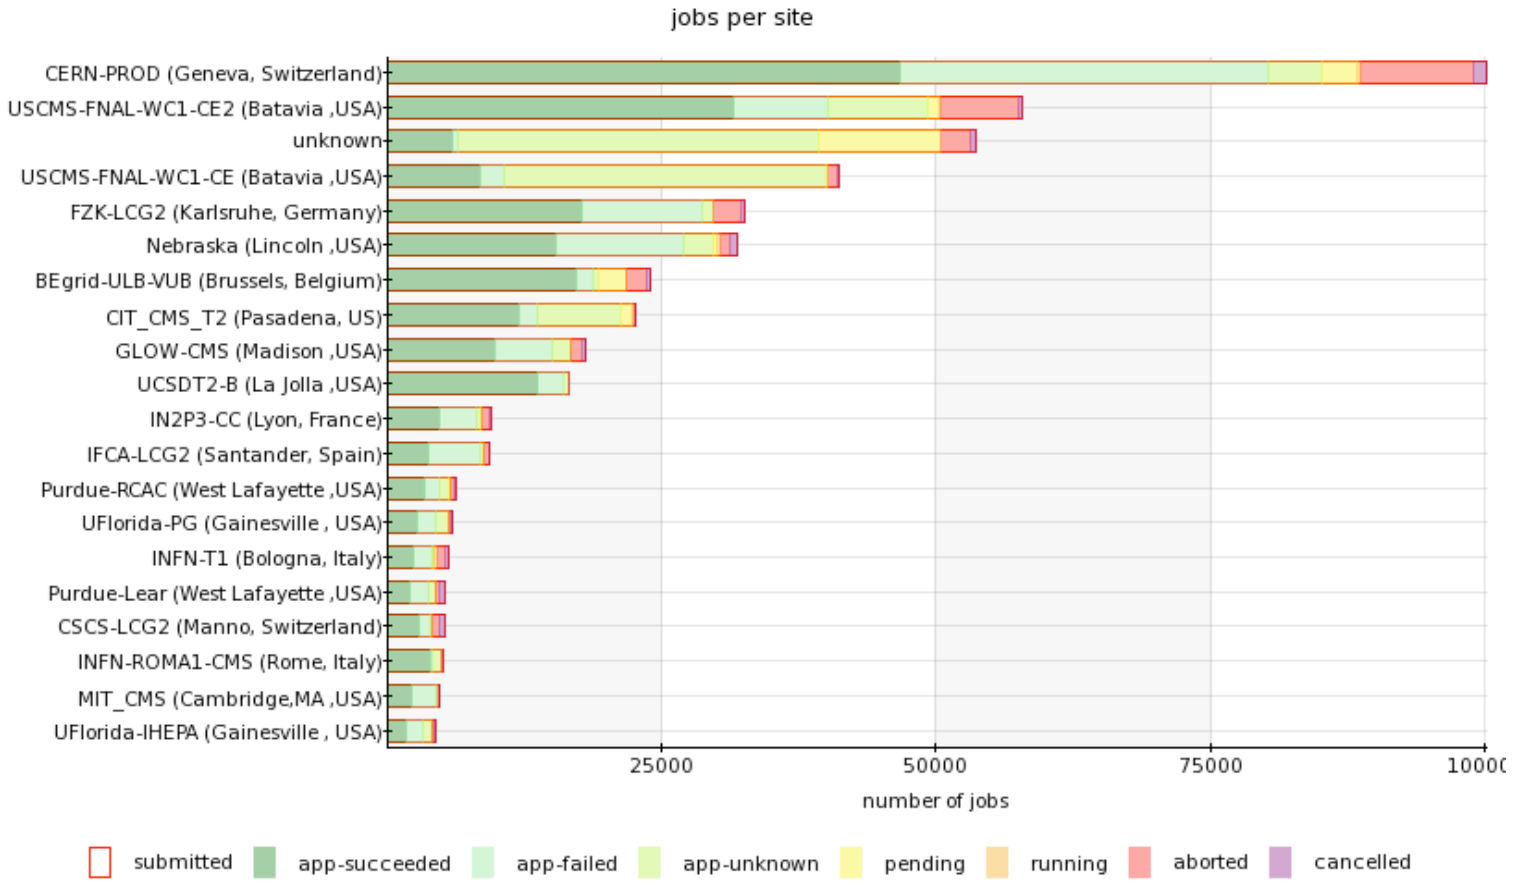

Restoration from Tape

Something that we haven't successfully done at any site is a demonstration of the ability to restore data from tape efficiently while simultaneously writing incoming files

- There is a wide range of opinions on what is the best way to do this
  - Some sites would prefer us to simply let the tape systems function and restore files as needed
    - This requires us to be smart about the data layout by families of tape
      - Worse when we need to read two files together
    - Potentially has a CPU efficiency cost to pay, but varies based on the performance of the tape system
  - Some sites would prefer organized pre-staging of data
    - Requires either some tools development, or verification that the SRM2.2 tools scale sufficiently to do this work