Reconstruction

EVENT BUILDING

- Nothing in online requires the assembling of **all** events (we do not employ level2 triggers, all events selected by our level1 triggers are recorded)
- Moving the event builder to the offline world makes it a lot simpler
  - The offline event builder does not have to keep up with peak rates
  - In offline we have many more cpus at our disposal
  - Crashes can be easily debugged
  - No loss of data due to event builder issues
  - Subevents are ordered in raw data files
- Disadvantage: Need to deal with ~60 input files in data reconstruction
- Still need to build a fraction of the events for monitoring purposes
- Combining triggered with streamed readout is going to be fun

Reconstruction

Highest priority: Save data into mass storage (hpss), hpss cache might pose I/O limit
- Provide enough capacity to copying data twice from counting house to rcf
- chances are that data in hpss only becomes available after the run ends

2 weeks for initial, one time calibrations determines currently the buffer disk space

15 secs reconstruction time: 90000 cores to keep up with incoming data
- Every second translates to 6000 cores
- Need to push 10GB/s through the whole system
- Need to combine data from multiple input files

Two passes envisioned – the first pass needs to take care of the heavy cpu lifting
- Save enough information on output to apply finer (“final”?) calibrations
- Analysis should be possible using the first pass output

Reconstruction and Analysis

Data objects not defined yet, any recommendations?

We are probably willing to pay cpu penalty for smaller filesize

Probably we need to save versions with selected events for selected fast analysis topics

At what point can we filter off those interesting events? During 1st pass already?

Possibility to analyze some topics without the TPC (based on Si tracking)?

Bulk of data will be analyzed using a train approach, especially if this requires reading from tape

Reconstruction + Analysis Flow (simple view)

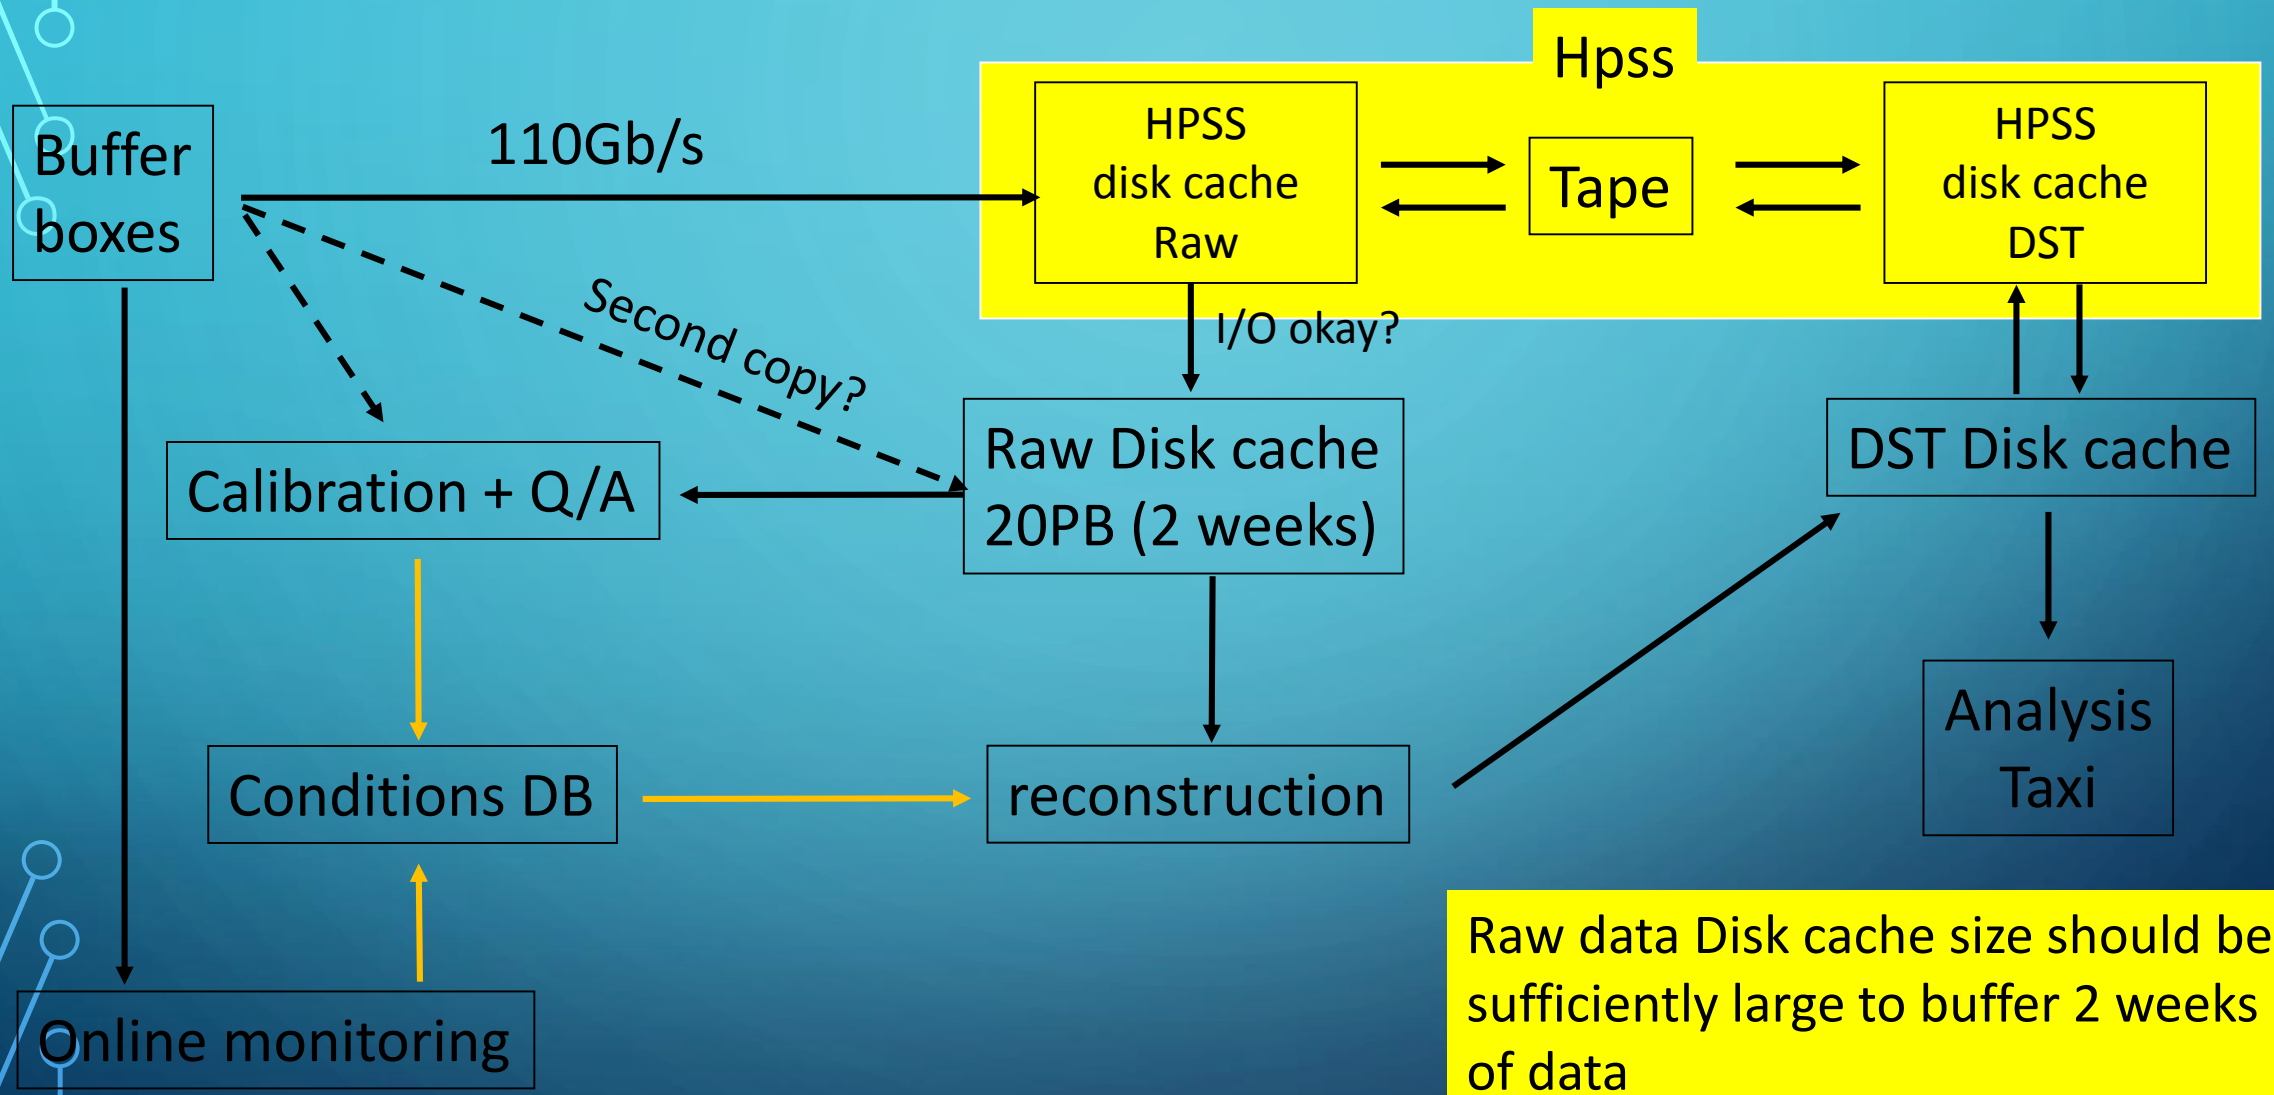

Raw data Disk cache size should be sufficiently large to buffer 2 weeks of data

Reconstructed data disk cache should ideally keep all reconstructed output