

BIG DATA IN HIGH ENERGY PHYSICS

Igor Mandrichenko

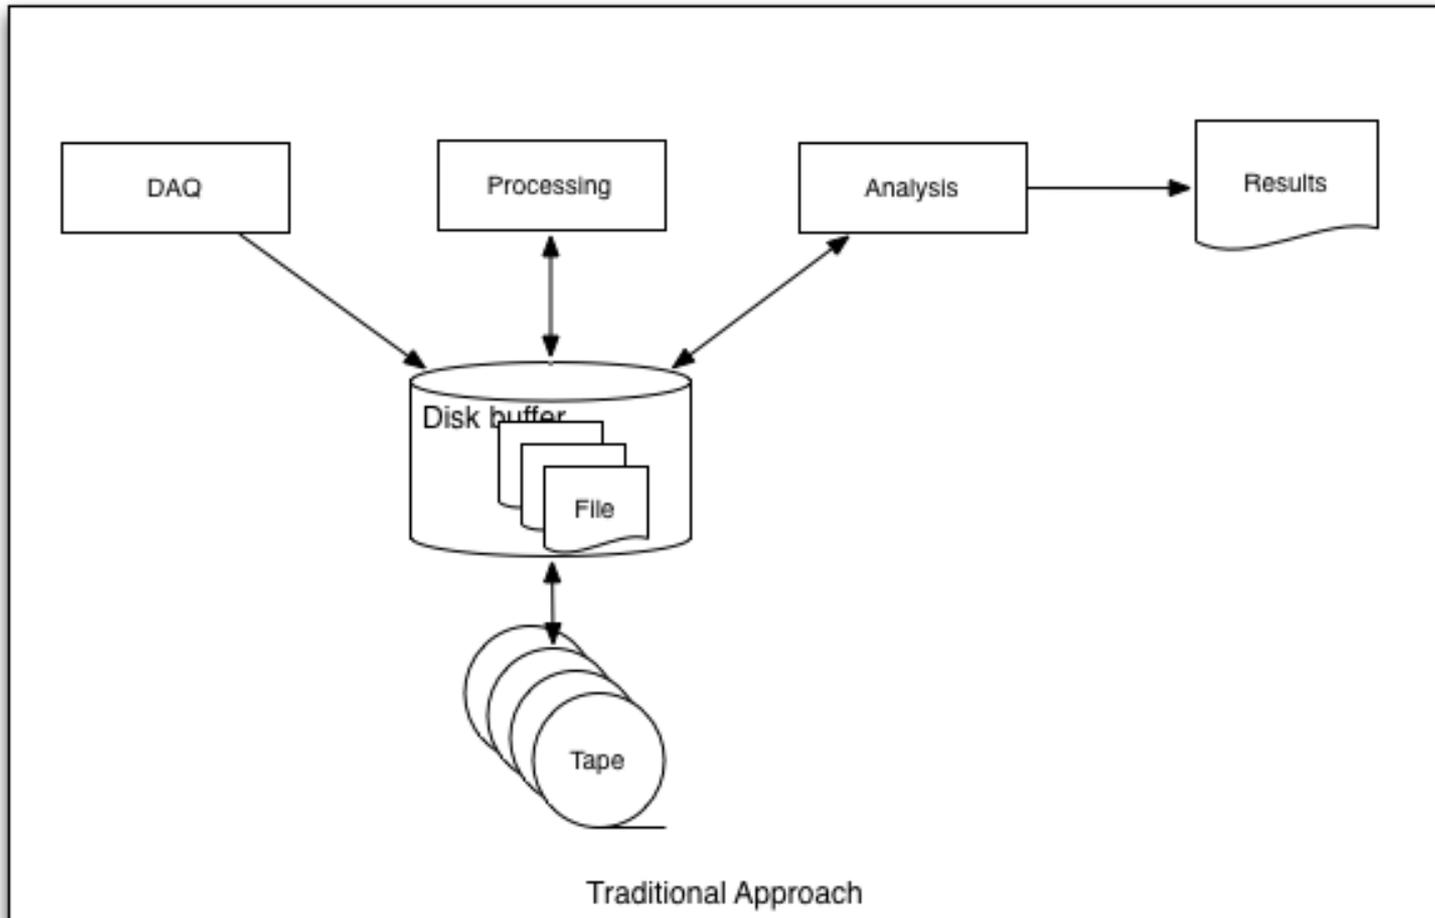
What is Big Data

- Big Data is *not* about size
 - It has been known for decades how to collect and record petabytes of data
- Big Data is about the ability to quickly analyze large amounts of data
- Data can be collected and stored quickly, and then always “immediately” available for the processing (reduction) and analysis

Traditional HEP approach

- Collect data (DAQ)
- Write raw data to tape as fast as you can, via disk buffer
- Process or reduce data (reconstruct events)
 - Read data from tape to disk
 - Write reduced data to disk and then back to tape
- Analyze data
 - Read reduced data (and often raw data) from tape to disk
 - Filter “interesting” data – read each event, decide whether it is interesting
 - Save “interesting” data for future analysis (write to disk and then to tape)
 - Every group or individual saves their own “interesting” sets, duplicating data
 - Further reduce data into physical result (mass, cross-section, etc.)

Traditional HEP approach



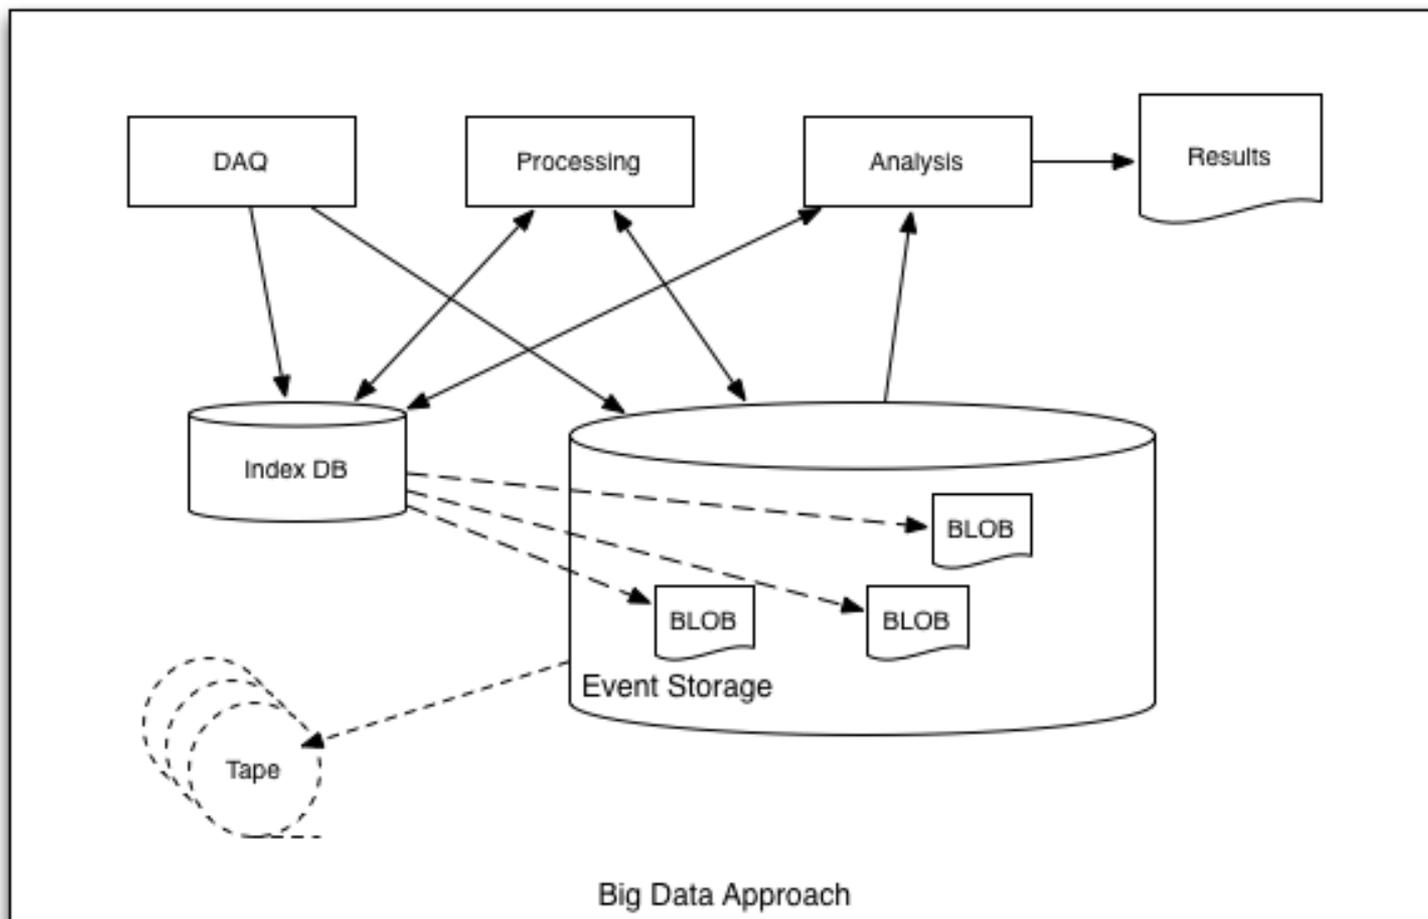
Traditional HEP approach

- Tape is primary storage of all data from raw to intermediate results of the analysis
- Disk is a buffer through which tape storage is scanned. Not considered reliable
- Data need to travel between tape and disk many times
- It is important to be able to filter and save “interesting” data once and reuse it many times so that there is no need to re-scan whole dataset again and again

Big Data Approach

- Collect data, store on disk
 - Index data
- Data processing (event reconstruction)
 - Read new raw data from disk – use index to define what is new
 - Reconstruct
 - Write reconstructed data to disk (associate with raw data)
 - Index data
- Analysis
 - Use index to find potentially interesting events
 - Create and populate your own index
 - Analyze interesting data
 - Update your index
 - Produce physical results

Big Data Approach



Big Data Approach - differences

- Disk is primary storage of data
 - 3-5 times replication, multiple physical locations
 - Can save data to tape
 - Probably never have to read data from tape
 - Representation – open, can be anything
- All data is always immediately available from disk
 - Direct access as opposed to “sequential” in case of tape
- Indexing instead of filtering and copying
 - Each piece of data exists in one (x number of replicas) instance, usable by anyone
 - Group or user can create arbitrary indexes

Big Data Approach - details

- Event Storage
 - Key -> BLOB database or storage, no SQL database
 - Event ID -> event data
 - Disk based, distributed, replicated, ideally, over WAN
- Index Database
 - User defined criteria -> event ID
 - Can be a relational database
- Indexing process
 - Map/reduce approach seems to fit well
 - Periodically running jobs, which analyze event data and populate indexes