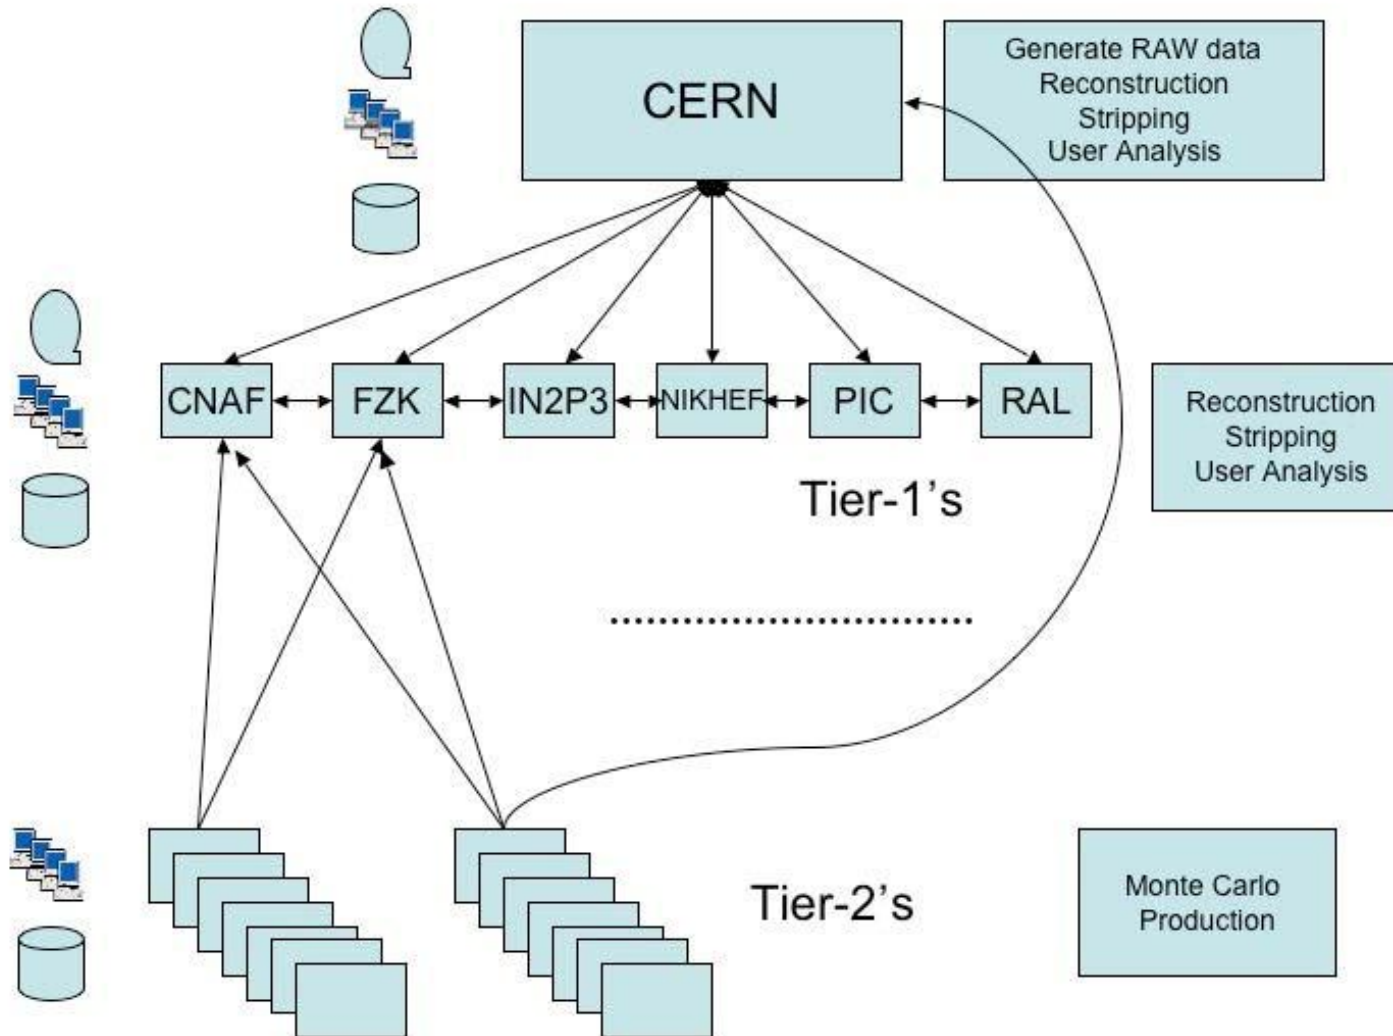




LHCb Data management in the future

Raja Nandakumar
(with Andrew C. Smith)

Current computing model



Current model and goals

➤ Current model

- ▣ DSTs from real data replicated to 7 Tier1s

- ☀ 2 T1D1 and 5 T0D1

- ☀ Analysis jobs candidates to run at all Tier1s

- ▣ DSTs from MC replicated to 3 Tier1s

- ☀ 2 T1D1 and 1T0D1

- ▣ This is the same regardless of the usage patterns of the files.

➤ Goal is to make more efficient use of disk

- ▣ More copies of popular data

- ☀ Up to 7 copies at Tier1s

- ▣ Fewer copies of “stale” data

- ☀ Reduce to only 2 T1D1 copies

- ☀ Eventually migrate T1D1 to T1D0 (envisaged in current model)



Outlook

- Many developments that interest LHCb
 - ▣ Nice to include in the middleware, stuff that LHCb has had to implement on their behalf
 - ☀ Re : Philippe Charpentier at the data management Jamboree and Ph.C. & ACS at Imperial
 - ▣ Note : Data management is tightly coupled to workload management
 - ☀ Need for pre-staging data as needed

Interesting stuff contd ...

➤ Improvement of catalog

- Both hierarchical & flat space logical names
 - ☀ Better URL construction, possibly splitting the name and saving the parts (protocol, SEName, SAPath, ...)
- Datasets integrated into the catalog
- Ability to handle SE statuses smoothly
 - ☀ Hide replicas when SE is down
 - ⇒ Automatic notification of problematic files?
 - ☀ Easy SE redefinition when needed
- Consistency with storage
 - ☀ Unreachable files
- Quotas (DN / FQAN)

Interesting stuff 3 ...

- Replication of “hot” data alone out of CERN
 - ☀ Hot = data high in demand
 - ☀ Depends on ability to reliably identify hot data
 - ☀ Potential to reduce storage need at Tier-1s
 - ⇒ 10%?
 - ☀ Medium / long-term future
- Metrics to evaluate cost of replication?
 - ▒ Replicate vs Staging vs Disk copy?
- SE accounting
 - ▒ “du”-like ability
 - ▒ Currently very heavy to do this using LFC