



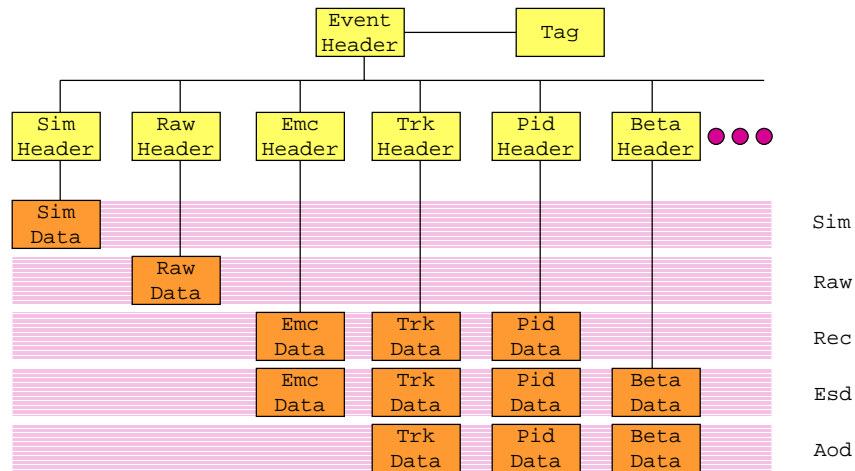
Data Distribution

David R. Quarrie
Lawrence Berkeley National Laboratory
DRQuarrie@LBL.GOV
for
Jean-Noel Albert
LAL-Orsay

Goals & Problems

- Goals
 - Bulk exchange of data between Regional Centers & Institutions
 - Exchange of event samples between physicists
 - ◆ Transfer of subset of data for event event
 - Must support bi-directional transfers
- Problems
 - Data from any one event is spread across multiple databases
 - ◆ In worst scenario, require 25GB to transfer all the data of a single event
 - Need to identify which databases correspond to the data sample
 - ◆ Logical to physical mapping
 - Database files have a unique identifier
 - ◆ Must avoid identifier collisions between source & destination

Event Structure & Placement



5/25/98

David R. Quarrie: Issues

Bulk Transfer Concept

- Database ID preallocation
 - Range of database Ids allocated for each site
- Creation of new databases constrained within allocation
- Ensure that all information from each event doesn't span database files
 - Within each tree
- Identify the database files at source
- Ship the files to destination & attach them to federation
 - Shipping via tape or network
- Provide navigational access to new data

5/25/98

David R. Quarrie: Issues

Physicist Data Exchange Concept

- Database preallocation
 - Range of database Ids for each physicist
- Loop over required events in input sample
- Create transient events
- Re-output required information to a single *Exchange* database (or set thereof)
- Ship exchange database(s) to destination & attach them to federation
- Provide navigational access to new data

5/25/98

David R. Quarrie: Issues

Status

- Database allocation data model implemented
- Application to assign allocation to sites & users not yet designed or implemented
- Database creation within allocation works
- Application exists to identify databases in an event collection
- Objectivity utilities exist to extract database files & attach them at destination
- Mechanism to provide navigational access to new data doesn't yet exist

5/25/98

David R. Quarrie: Issues

Future

- Possibility of using Objectivity FTO/DRO for distribution
 - Fault Tolerant Option
 - Data Replication Option
- Multiple *images* or replicas of database at different sites
- Acts as distributed database cache
- Read access from “closest” image
- Updates causes synchronous updates at transaction commit time to all images

5/25/98

David R. Quarrie: Issues

FTO/DRO Problems

- All transfers via network
 - Cannot export/import databases to/from tape
 - Database bandwidth & reliability is limiting factor
- Some operations require all sites to be active
 - Creation of a new image
 - Schema updates
 - A remote site could hold SLAC operations hostage
 - Due to be fixed in next Objectivity release
- Synchronous update commits could degrade SLAC
 - Only create images after “real-time” databases created & stable

5/25/98

David R. Quarrie: Issues


FTO/DRO Problems (2)

- No concept of “closeness” of remote images if no image available locally
 - Cannot adjust image topology on a per-site basis
 - ◆ e.g. LBNL & Annecy might view IN2P3 as nearest neighbour
- Objectivity are aware of the problem and we are discussing possible solutions with them
- Timescale unclear

5/25/98

David R. Quarrie: Issues

Data Distribution Strategy



5/25/98

David R. Quarrie: Issues