

# Minerva Analysis ntuple production

Philip Rodrigues

April 8 2014

Minerva data management review

# Overview: AnaTuple production stage

- In: reco output file
  - POOL format, complex, same for all analyses
- Out: Analysis ntuple
  - flat ROOT tree, simple, one type of ntuple per analysis
- Some custom reco done, some simple putting data into branches

# Usage pattern

- ~20 analyzers each running their own analysis
  - Inputs same for everyone, output varies quite widely
- MC dominates data (in CPU and file size)
- One analysis pass reads in  $O(10k)$  files of ~700 MB each. Total ~20 TB
- Each analyzer does this ~every few months? Smaller test processings in-between
  - Crunch times before conferences see higher usage

# Each job

- Each job reads several (5-10) input files, makes one output AnaTuple
  - So that's ~5GB input (MC)
  - Output file (MC) can be up to ~200 MB, often smaller
  - Minor outputs: logfiles (plain text), options file (plain text), histograms (root file).
- Input can come from bluarc or pnfs.
  - Optional: not all users using pnfs yet. Will be default soon
- Output all goes to bluarc, not catalogued.
- Time taken varies by analysis, but ~1 hour per job
  - (This is processing time, not including waiting for locks etc)
- Number of input files merged is definitely negotiable

# Past/current problems, areas for improvement

- Bluarc locks leading to low efficiency, long wall times
- DB server unresponsive (ana stage accesses DB for “kludge tools”)
- Diagnosing problems when they occur. Who to call?
- Understanding our grid allocation and making sure we're getting it
- Submitting a large pass does a lot of work on the interactive nodes. Can this work happen on the grid nodes? Will make interactive machines happier
- Can we write output to pnfs and read it back in ROOT jobs? How?