

Rob Snihur

Simulation mtg., Mar. 10, 2005.

# Concatenation status

- Old way
- New way
- Features
- Error recovery
- Testing
- Output filesizes
- Errors
- Plans

# Concatenation

- Old way (Summer2004)
  - Uses 2 tape loggers: fcdldata004 & fcdldata031
  - Uses 3 disk buffers: fcdldata004, 005, 123
  - 2 scripts:
    - Make\_prod\_concat\_tcl (then update tarball)
      - interactive
      - Uses “cafhostdir” to read files from disk
    - Submit\_prod\_concat
      - Runs on CAF in high-priority queue
      - Use stntuple\_maker.exe
      - Access to input (& output) files via rootd
  - Human must manage resources, i.e. disk

# New proposal

- One DCAF segment: simulation makes many (small) output files, totaling X GB (where  $X > 1$ )
- At the end, concatenate all small files within the DCAF segment

MCful I :

```
For (i = i start; i < i ed; i ++ ) {  
    Call MCPProd;  
    cp $outfi l e $outdi r  
}  
Concatenate $outdi r
```

# Features

- Simple
- self-contained
- No need for durable cache
  - instead use disk available on worker nodes: **how much ?**
- No need for separate submission of concatenation jobs
- Less book-keeping
- **Average user can do it**
- Fewer segments to submit

# Drawbacks

- Need to know event size for each sample in advance
  - Could run a 100-event test job

# Error recovery w/in CAF job

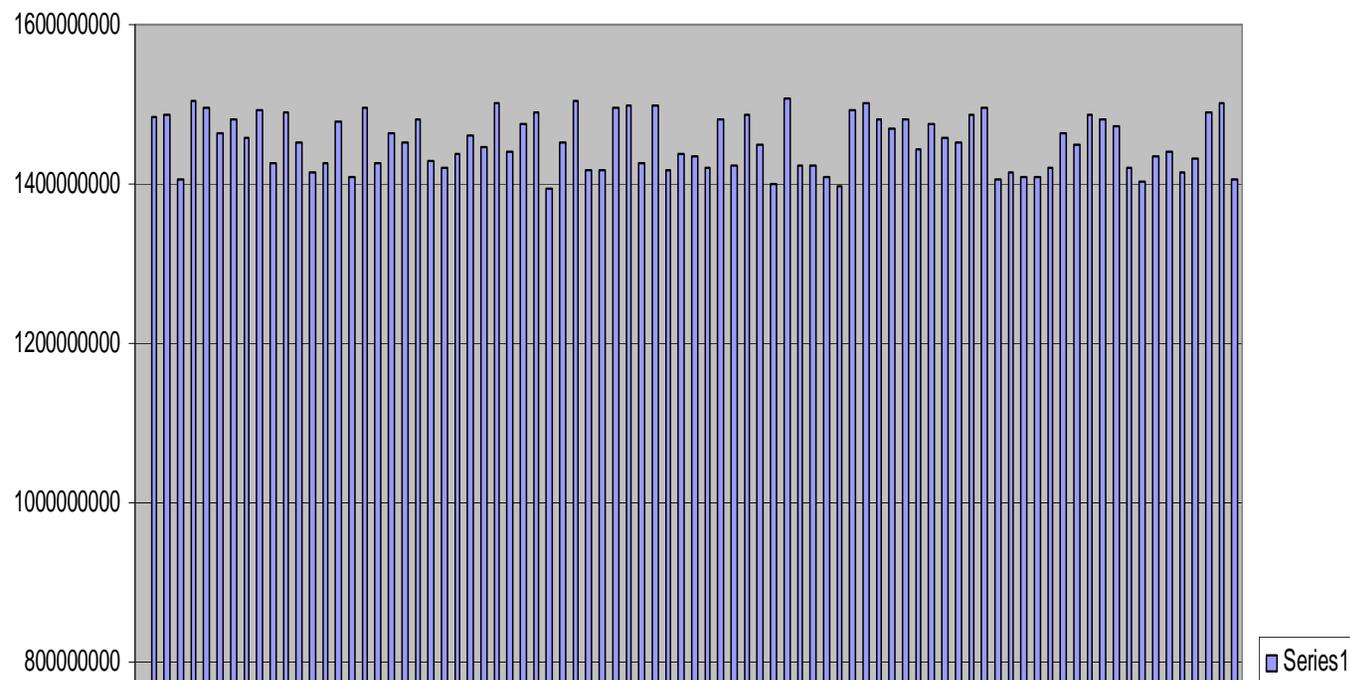
1. Simulation (MCProd) failure
  - a) Retry MCProd (max\_retry times)
  - b) If repeated failures: Send email, Ignore it & proceed
2. Concatenation failure.
  - Check concatenation log within CAF job
  - Transfer output file to tape logger node
    - If transfer fails, send email & retry 3 times
  - Run Edm\_EventLister on remote file
    - Verify # events matches that in concat log
  - Email if any problems

# Testing

- Generate dummy sample: ~1 M t-tbar events
  - 3k run sections (~100 MB / section = 430 events / section)
  - 176 segments (~1500 MB / segment = 6k events / segment)
- Run at SDSCCAF (2.4 GHz CPU)
  - Estimate 22h of real time / segment
    - 80 completed
    - 68 running
    - 28 pending

# Output Filesizes

- Target was 1.5 GB
- Observe 0.1 GB variations, as expected



# Errors

- 20 MCProd failures (each failed twice)
- Many (79 so far) concat failures
  - Some segments need to run 2 concat jobs (reason unknown)
  - Only 1st concat job succeeds
    - Reason known
    - Fixed in earlier version
    - Forgot to include fix in new version
- File-naming problem
  - Theory: old bug returns; fix not included in new version
  - Duplicate event numbers in last runsection
  - May affect output file sizes

# Plans

- Pick up 2 missing fixes
- Add timing info
- Port to latest cdfsoft2 version
- Rerun stress tests
- Replace DFC tape logger with SamStore
- Full tests (write dummy datasets to permanent storage?)