Status of OO Production in the US

Hans Wenzel Fermilab
April 11 jetsmet meeting

- Update of Status of OO Production in the US and at Fermi

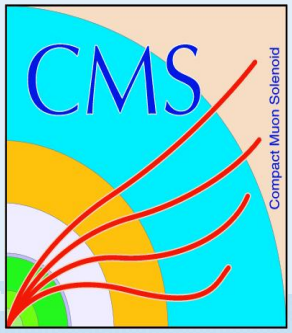

Status of OO Production at Fermilab

- FNAL production page, links to tools, data, scripts and scripts
  http://computing.fnal.gov/cms/Monitor/cms_production.html
- I am processing a list of all the data availabledata sets.
  http:/home.fnal.gov/~wenzel/data.html (wip)

Status in the US

- We started wekly meetings with UCSD, Caltech, U. of Wisconsin.
- UCSD, Caltech have imported minbias hit federation and data from CERN.
- Wisconsin importing hit federation.
- They think they can start production soon. Will start with signal data sets not currently at FNAL

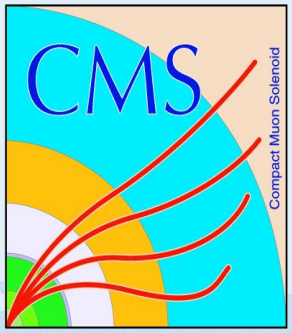

Hardware and software used in Production at FNAL

- **2 Dell Servers with 4 Pentium CPU's. GALLO is the IO node of the FARM and VELVEETA is the Objectivity Database host. Each has > 200 Gbyte raid attached.**
- **40 dual CPU (popcrn01−popcrn40). For Digis production 9 are used as pile up server leaving the other 29 (58 CPU's) as worker nodes. Currently we have a sample of 500 000 minimum bias events. @2 nodes for R&D**
- **STK robot (enstore)**
- **FBSNG batch system −−> can only recommend**

WORKERS:

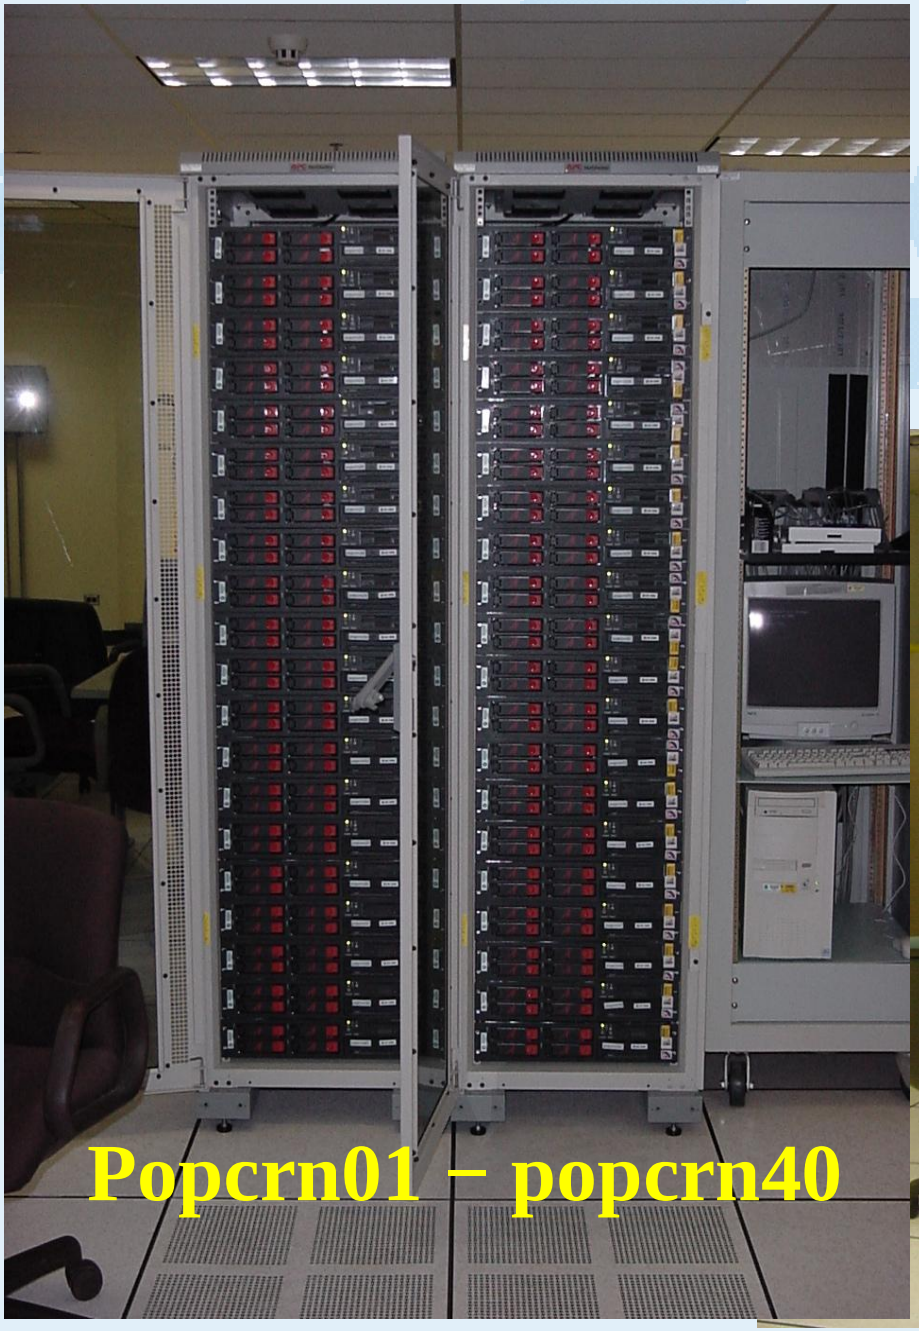

Popcrn01 – popcrn40

SERVERS:

GALLO,WONDER,VELVEETA, CMSUN1

OODigis

OOHit Data

GALLO

VELVEETA

Pile up

Popcr31

Popcrn32

Popcrn40

Popcrn01

Popcrn30

AMS Servers on
popcrn31−popcrn40
(pile up servers)
GALLO (data server)
VELVEETA (DB host)

Every one talks to everyone

Current Status

- We use ORCA 4_4_0_optimised
- Imported Pile up sample from CERN of approximately 500K events. Split over the 9 popcrn server nodes
- Created User federation to access set OO_FD_BOOT to cmsun1.fnal.gov::/usr/cmsun1/r5data08/UserFederation/jet0900.boot At the moment contains: jm_msugra1, jm_msugra2, jm_msugra3 no pileup and 1034. jm_hplus_m200 1034 we keep adding data. Pal and Pam ran ntuple makers on the Federation and say it's fine. But one caveat trigger

At FNAL

- We have started to run production like an experiment. We have daily operations meetings involve sys admins, mass storage group, batch people.
- Reconfigured the batch system to be more efficient: 2 nodes taken out for R&D, 9 pile up servers now can also run CPU intense "no pile up" jobs.

FNAL

- When will it all be over? Used data from last night. Back on the envelope 2.5 million events x2 (no pile up 10^34)
- 10^34: 29 nodes x2 CPU's x3jobs/day x500events = 87000 events/day
- Nopileup: 9 nodesx x 2 CPU's x8jobs/dayx500events = 72000 events/day
- Makes 159 000 digis events a day.

The big buts

- depending on load CPU efficiency varied between 46% and 86 %? For the estimate I used pessimistic number (50%)
- Stop digis production for OOHit.
- Stop farm for backups. Closed db files we can backup during running but but carf meta data only while no DB transactions are ongoing. We investigate if we can get around it.

When will it be done

- Say 120000 events/day might be a realistic number
- We have 2500000 jetsmet events. 5000000 if we require 10^34 pileup and nopileup. Just with the resources at fermilab it would take 42 days to process!!!

Current Status(cont.)

- We have tools in place to duplicate data bases, delete selected files et. (thanks to Shahzad Muzaffar) scripts are in CVS
- All scripts are in place to run production in CVS and tagged if some one wants to have a look at it.
- Scripts being ported to cern and are evaluated in the other US centers.