CAS meeting Feb. 20 2001

Hans Wenzel Fermilab

- What have I been doing
- Status of OO Production at Fermilab
- Plan for the next few weeks

Status of OO Production at Fermilab

- Plenty of production data in enstore tape robot: http://home.fnal.gov/~cmsprod/monitor/report.html
- Data cards: http://home.fnal.gov/~cmsprod/Production.html managed in CVS scripts to create batch jobs
- In addition some extra samples requested by different members of jetsmet group e.g. Z−prime with different masses, tth, single pions with eta cuts, minimum bias for pile up.
- Some work under way to establish a Production system which would work on different RC (Meetings with Wisconsin, Caltech and FNAL)

Status of OO Production(cont.)

- Production Hardware:
- 2 Dell Servers with 4 Pentium CPU's. GALLO is the IO node of the FARM and VELVEETA is the Objectivity Database host (Production Federation). Each has > 200 Gbyte raid attached.
- 40 dual CPU (popcrn01−popcrn40). For Digis production 10 will be used as pile up server leaving the other 30 (60 CPU's) as worker nodes. Currently we have a sample of 420 000 minimum bias events.
- CMSUN1 8CPU sun machine. Host of User Federation

OOHitFormatting (min bias)

- Massive data movement

Stage in FZ−files from Tape robot

↓

OOHit into Objectivity(VELVEETA)

↓

VALIDATE DB + CLEANUP

↓

SPLIT DATABASE
COPY TO PILE UP SERVERS

10 PROCESSES AT A TIME

BOTTENECKS:
NETWORK

OOHitFormatting (data)

- Stage in FZ−files from Tape robot
- OOHit into Objectivity(VELVEETA)
- VALIDATE DB + CLEANUP
- SPLIT DATABASE
  COPY TO GALLO

10 PROCESSES AT A TIME

BOTTENECKS:
NETWORK

OODigis

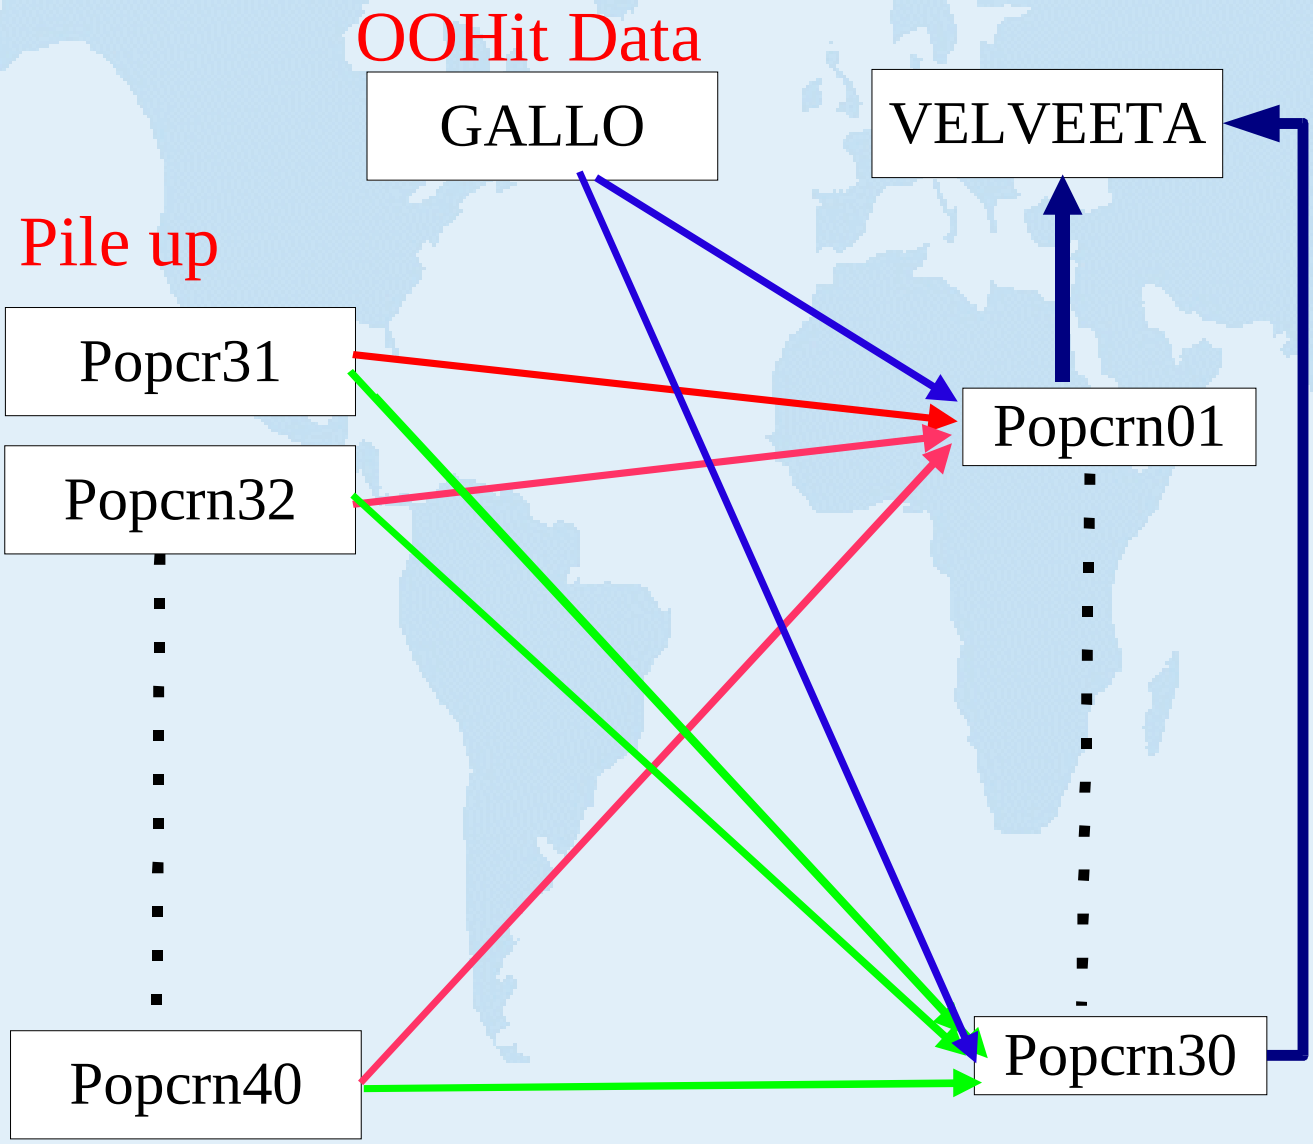

AMS Servers on
popcrn31−popcrn40
(pile up servers)
GALLO (data server)
VELVEETA (DB host)

Every one talks to everyone

Two weeks ago

- DB hosts and Pile up servers should have Gigabit ethernet connections (investigating we will run some tests this week CMSUN1 is now connected with Gbit ethernet). We will purchase some cards for the LINUX boxes and run some tests in the next weeks. Fortunately the switch has a few ports.

Two weeks ago

- Test run with 5 pile up servers (ORCA 4_3_2)
- 100 000 min bias events (had more but a failed tape drive killed 75 % of the sample)
- Monitor exactly what is going on network, database locks etc. Most of the info is on the www.
- Found we need more than 5 pile up servers
- Need to start AMS with file limits set higher (Linux 1024 not enough). Opening more threads doesn't help.

Current status of OO–production

- Upgraded to ORCA 4_4_0 (redo to get new version with optimization switched on)
- Updated the batch system (can be kerberized)
- Produced more pileup 500,000 (this is sufficient for signal samples but what about bgr.samples which have more than 150000 evts.)
- populated the min bias data base

- Backed up min bias sample now we are in the process of splitting the database over the 10 pile up servers
- This week we will run OODigis starting with jm_hplus_m200 and jm_hplus_m400 data sets

Plans for the coming weeks

- Do OO production! Develop tools to automate and monitor. Everything will be avalable on the web (network, cpu, db locks ....)
- During production the Objectivity database on velveeta shouldn’t be accessed by users. So we need to duplicate the DB and create a user federation (CMSUN1 as host?)
- Validate (run ntuple makers etc.)

Plans for the coming weeks(cont.)

- Next step would be to get involved in software development probably OSCAR. There is a Geant4 workshop in Ann Arbor Michigan starting Feb. 26. I will check if I can attend. Hopefully my C++ skills didn't evaporate.
- March 5−18 I will be at CERN for CMS week and the week after. Hopefuly it will be a good time to discuss producton and OSCAR involvment.