Template for Neutrino Experiment Computing Infrastructure

The Mu2e Experiment

Rob Kutschke, FNAL CD

Caveat

- Some of these numbers are official collaboration numbers taken from the proposal.
- Many are just my guesses, not vetted by my colleagues.

Experiment specifics

- Purpose: to search for the coherent, neutrino-less conversion of a muon to an electron in the electric field of an atomic nucleus.

- Number of users: O(75)
  - Remote: O(60)
  - Using Fermilab facilities: O(30)
  - Based on O(150) total collaborators; today O(75).

Experiment schedule

	'09	'10	'11	'12	'13	'14	'15	'16	'17	'18	'19
Planning	XX	XX	XX	XX	XX						
Construction				X	XX	XX	XX	XX			
Commissioning									XX		
Data taking									XX	XX	
Data analysis									XX	XX	XX

Working plan is:
CD0 – about now
CD1 – Oct 2010
CD2 – Oct 2011
CD3 – Oct 2012
CD4 – Oct 2016

Official plan is 2 years of commissioning + data taking. I believe that we will need more like 3-4 years.

Data

- How many events/year?
  - Test beam << experimental data.
  - Pedestal and calibration – $2x10^9$ (alternate running configs)
  - Normal data – $2x10^9$
  - Normal data after quality filtering? – WAG $O(10^7)$
- How large is each event?
  - Zero suppressed – O(50 kB)
  - Non zero-suppressed – (600 kB)
  - Reconstructed – WAG O(5 kB); 1 track + subset of raw.
  - Simulated – WAG: reco + 20% = 6kB
  - Data summary – WAG: = reco = O(5kB).

Central FNAL systems

- CPU used (see table)
- Storage used (see table)
- Uses:
  - Reconstruction and data filtering
  - Calibration and alignment
  - MC Generation
  - User data analysis

Data flow

	Pre-2015	2015	2016	2017	2018	2019	2020
Raw Data, TB				200	200		
Processed Data, TB				20	20	20	20
User data, TB			Small	Small	Small	Small	Small
Simulated data, TB	60?	120?	240?	240	240	240	240

WAG: MC events = 10 x reco data.
WAG: reprocess the full data a few times after run
WAG: MC needs decay by ½ in years before startup.

CPU needs

	Pre-2015	2015	2016	2017	2018	2019	2020
Running							
Reconstruction				100	100	100	100
Calibration				??	??		
Skimming				Small	Small		
Analysis				Small	Small	Small	Small
Simulation	200?	375?	750?	750	750	750	750

Please use CPU-years on a current machine
e.g. # events * time per event in sec * 3 x10^7 * reprocessing factor

- All WAGS: see next 2 pages
- Includes resources at FNAL and remote.
- Scale: 1000 cores DC; MC dominated.

Reco CPU

- L2 Trigger farm has 200 cores.
  - Has 5x headroom (for noisy data).
  - Does 5x data reduction.
  - Is only used 2.E7 seconds per year.
- WAG: Offline reco is 10x more than L2 trigger.
  - Need 50 cores
  - Reprocessing factor of ≈2 implies O(100) cores.

MC CPU

- WAG:
  - MC reco same speed as data reco = 50 cores
    - Not 100 cores since we do this once.
  - MC gen+sim = 50% of reco = 25 cores
    - Since sim overlays noise (fast) that slows reco down.
- We need something like 10 MC events generated for every reco data event:
  - So (50+25)*10 = 750 cores.
- For years before startup: assume factor of 2 less each earlier year.

Operating systems

- What OS is used?
  - We are developing for SLF, late SLF4 and higher.
  - We might consider MacOS X.
  - Do not intend to support Windows.
- Do all collaborators have to use the same one?
  - They may use any of the supported platforms.
  - They are on their own otherwise but root based analysis will likely work.

Data storage and tracking

- How do you catalog data?
  - Acquire a file catalog system: GRID friendly!
  - Lots of blue arc or equivalent. dCache?
  - Guess: tape is needed for archival/backup only.
- How do you provide remote access to data?
  - Not yet known. Options are:
    - ftp from FNAL blue arc.
    - Whatever grid service is both recommended and supported by CD at the time that we need it. Let others be trail blazers.

Remote systems

- How many remote institutions provide resources for your users/collaboration
  - Don't know. Probably significant offsite MC.
- Do they have special systems for you or shared? Not yet known. Likely both.
- What is done at remote institutions?
  - Reconstruction – I would be very surprised if yes.
  - MC generation – Likely a lot.
  - User analysis – Yes

Data distribution to remote sites

- Where are data distributed:
  - To most collaborating institutions
- What kind of data:
  - Skimmed summary data plus skimmed MC?
- How much data: O(30) TB/year = 1 MB/s
- How fast does it need to move
  - Burst rate 10x DC rate? 100x?
  - No idea what is realistic?
- What method is used: both push and pull.

Grid

- Do you use the Grid? Not yet but we will.
- Do you use Grid tools such as Gridftp?
  - Expect to later. Initially FermiGrid + blue arc.
- Do you use Glide-in or some other tool?
  - Don't know yet. Will follow recommendations.
- Do you use the FNAL Grid exclusively or do you use more general grid resources?
  - Will develop for vanilla grid use and avoid local dependencies. But initially a blue-arc dependence?

Databases

- Technology used? Whatever CD will support.
- Size: Don't know.
- Access rate: Don't know.
- Are they replicated remotely?
  - Snapshots yes. Need to be able to run on a laptop that is on an airplane.
- What is stored: conditions data, file catalog, data quality monitoring summaries, …

Conditions

- How are conditions and calibrations stored?
  - They will live in databases.
- How are they accessed?
  - Whatever protocols CD experts recommend and support.
  - Expect that snapshots of subsets will reside locally.

Code management

- Code repository
  - CVS – default. Would switch to svn iff CD will support an svn server and iff the build system plays nice with svn.
  - svn: has automagic, project wide versioning; simplified management of branches (ask Jim K for details).
- Build system
  - SoftRelTools – default choice but not yet used.
  - CMT – not familiar with it.
  - Now using scons as a gmake replacement. Good enough for the short term; for the long term?

Standard packages

- What standard packages are used:
  - GEANT4 - Yes
  - ROOT - Yes
  - GENIE - No
  - NEUGEN – No
  - LCG – We want to be able to use generic grids.
    - So probably.
  - CLHEP - Yes
  - PYTHIA – No
  - Infrastructure software: see next 2 pages …

Infrastructure Software

- Infrastructure:
  - Framework proper, services, persistency, conditions, configuration, build & release management.
  - Will use it as an external product, like G4, CLHEP …
- Mu2e requires something that is supported.
- CD's recommendation (Jim Kowalkowski's group).
  - Stripped down CMS infrastructure (mature).
  - Alpha release delivered in January:
    - Less DB and file catalog; persistency not full featured.
  - They are offering this as a new standard package.
- Use it from first non-real time element in DAQ/trigger …

Infrastructure Software II

- Received alpha release in January
- Learning how to use it wisely (toy detector):
  - So far just O(50%) of me for 8 weeks.
  - Iterate design as I add features to toy detector.
  - Integrated with G4.
- First release of Mu2e detector in next weeks.

What worked really well?

- Ask me again in a year

What would you not do again?

- Ask me again in 6 months.

A Final Comment

- We make lots of neutrinos.
  - We just don't plan to do anything with them!
- Is it a good thing or a bad thing if some of our neutrinos end up in your detectors?
  - What do you need to know from us to answer this? Energy spectrum? Rate?