Minos+ FY14/15 Computing Needs

Arthur Kreymer,Rashid Mehdiyev,
Robert Plunkett, Jennifer Thomas
Fermilab

Scientific Computing Portfolio Review
January 30, 2014

Scientific Goals for the Coming Year

Minos will continue analysis of current low energy and cosmic data.

Minos+ taking NOvA medium energy beam data from Sep 2013, to continue through 2016, probing :

- Sterile neutrinos
- Non-standard interactions
- Improvement of neutrino oscillation parameters
- Extra dimensions
- Time of Flight

Large Scale or out of band computing needed to complete these goals

Ongoing Minos analysis is predictable

- One major reprocessing per year

Minos+ Far detector beam/cosmic and Near cosmic data rates remain small

Minos+ Near Detector beam data rates will be 9 times Minos

- Increased POT and neutrino flux
- We assume 6E20 NovA beam for all FY14
- Net ND data rate increases by 3.6 , including cosmics

Minerva will require reconstruction of the full ND data set

- Still fairly small, of order 1000 core-weeks on Fermigrid

Did we meet FY12 Scientific Goals ?

The Minos FY13 goals were met, using the budgeted resources.

Opportunistic Fermigrid resources played a major role

- See details on the final Resource Model slide

Data Taking Needs

FEF maintains the Control Room and DAQ computers

Minos shift activities use tools including

- ECL
- Redmine
- Mysql conditions database
- IFBEAM beam data logging

DAQ systems uses

- Equipment from PREP
- Networking support by NVS
- Data logging via Dcache and Enstore supported by CCD/DMS

Many shifts will be remote using the methods proven in FY12

Computing Services

Database

- DBA support for Mysql offline conditions database
- Size is 100 GB, expect growth under 20 GB/year

Storage systems

- Dcache/Enstore
- Bluearc
- AFS (home and web areas, legacy files)
- SAM data catalog

The Minos framework uses Root.SoftRelTools, UPS/UPD, CVS code repository

Minos grid jobs use Jobsub submission to HTCondor and GlideinWMS

Minos uses collaborative tools DocDB and Readytalk

Minos has a modest interactive cluster, and a few dedicated data handling hosts. Most servers are shared with other IF experiments.

Production

Timely data reconstruction is done on Fermigrid.

"Keepup" reconstruction is run daily on all data.

Raw data occupies 12 TB in Enstore and Dcache

- We keep all raw data on Dcache disk, to reduce tape risk
- We write a 'vault' copy of raw data to a separate tape robot

Minos+ will produce a few Tbytes/year of beam timing data

Reconstructed data uses 300 TB tape, 45 TB bluearc

- A single pass on all Minos data uses 20 TB tape

- SAM data catalog manages all production files

Analysis

Analysis uses the same Fermigrid resources as Production

Analysis uses about 75 of 220 Tbytes of Bluearc project disk

Minos uses a Mysql conditions database via root 'Tsql' calls

The modest interactive pool is well used

- Four 8-core systems, presently real, could become virtual
- Provides a local condor batch pool

The software framework is mature and stable

Peak analysis needs are usually before major conferences

- Moriond/La Thuile 2014 – Mar
- Neutrino 2014 –Jun
- ICHEP 2014 – Jul

Simulation

Minos Simulation is generated off-site at collaborating institutions, and imported for reconstruction at Fermilab.

Simulation uses about 400 TB of tape and 45 TB of project disk.

Collaboration Tools

Minos relies on DocDB for document management and to organize all meetings

Minos uses CS central servers for its primary web pages

Minos makes heavy use of Redmine

Minos general network needs are relatively modest.

> NVS supports networking at the near and far detectors, and the tunnel to Soudan via ESNET and UMN resources.

Minos uses Readytalk for almost all its meetings.

Minos makes good use of consulting from all of the CS

Minos is too small to deploy unique solutions, so carefully tracks and follows the direction of the Fermilab Computing Sector.

Offline Computing Coordinator – Rashid Mehdiyev

Online Support – Bill Badgett, Donatella Torretta

Resource Summary

NFS (Bluarc) Project disk usage will remain at 250 TB through 2016.

Dcache use will increase dramatically to include all managed datasets

Tape is the historic 100 TB/year in FY13, 90 TB in FY14,

- releasing 400 TB of unuseful candidate files in FY14

Grid slots shown are for Minos on GPGrid – modest increases each year

Interactive cores are the existing minos50-53

Most servers are shared with other IF experiments, and virtualized

- Retain two dedicated I/O servers long term

Resource	2011	2012 (+inc)	2013 (+inc)	2014 (+inc)	2015 (+inc)	2016 (+inc)
Disk NFS TB	160	240 (+80)	240	240	240	
Disk Dcache TB	46	46	46	250 (+204)	450 (200)	700 (+250)
Tape TB	650	750 (+100)	850 (+100)	550 (-300)	640 (+90)	730 (+90)
Grid Slots	1000	1200(+200)	1200	1600	2000	2500
Interactive Cores	32	32	32	32	32	32
Servers	6	4	2	2	2	2

Resource Model

Extrapolation model

- Assume all Tape and 80% Disk usage scales with ND data rate (conservative)
- Disk can be used to 80% capacity with good performance, so allow 20% overhead
- Calculate ND data rate increased 3.6 times from Minos
  - Assume beam data rate increase of 9x over Minos
  - Minos ND data is 1/3 from beam, 2/3 from cosmics

NFS (Bluearc) disk usage has grown about 50 TB/year historically

- This halted in FY13, with deployment of 4 PB of Dcache disk for managed datasets
- Bluearc is still essential as Project disk, for dynamic analysis

TAPE usage has grown 100 TB/year historically

- This continues through FY13 at 90 TB/year due to no writing 'candidate' files
- We will remove 400 TB of existing 'candidate' files in FY`14

CPU increased needs are assumed to be provided mainly opportunistically

- FY12/13 allocation was 1200 slots
- Average FY12 usage was 1000 slots, from Gratia accounting and local monitoring
- Peak usage was typically 2000, with a record of over 6000 slots