Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

ELBNF and 35-ton FY15 and FY16 Computing Needs

Thomas Junk, Eileen Berman

Scientific Computing Portfolio Management Team (SC-PMT) Review

4 March 2015

Important Dates for ELBNF, FY15 and FY16

- Jan 22 – 23: ELBNF Proto-Collaboration Meeting
- May 28 – July 28: 2015 35-ton Phase 2 Operation
- July 2015: CD-1 Refresh: Need to Update CDR
- 2017 --? CERN beam test
- 2021 – time early for FD installation start

- Future dates (more or less in order)
  - Spokesperson Elections
  - Formation of New Collaboration Working Groups
  - CD-3a
  - CD-2
  - CD-3

Scientific Goals for FY15 and FY16

Optimize Neutrino Beamline Design

- Target, Horns, Decay Pipe, Absorber, Hadron Monitors
  - Predict physics performance of different high-power designs (1.2 MW and above)
- Other components – target chase, shielding, baffle, magnets, embankment are designed by AD (LBNF Facility)
- Target, Horns, etc. also designed by the LBNF Facility, but in conjunction with ELBNF requirements, optimization

  They must be buildable, operable, safe, and reliable.

  And we want the most physics out of the beam
  - Mass Hierarchy
  - δ_{CP}
  - Other physics

Scientific Goals for FY15 and FY16

- Assemble, Commission, Operate 35-ton Prototype
  - Two months of operation, May – July 2015
  - Extract performance parameters of the modular TPC design
    - LAr Purity
    - Signal and Noise characterization
    - Tracking Efficiency
    - Measure response to particles in the gaps
    - Measure space charge effects
    - Stress-test reconstruction of ambiguous data (wrapped wires)
  - Measure photon detector performance – 3 designs
  - Analysis of data and publication will take multiple years

Scientific Goals for FY15 and FY16

- Optimize Design of Far Detector
  - CD-1 Refresh. Summarize work done to characterize performance and optimize the modular 1-phase LArTPC
  - Many Ideas to study
    - Cavern shape – how does this affect acceptance? Can it be made to accommodate a 1-phase and a 2-phase detector so we can postpone the decision?
    - Wire angle optimization
    - Wire pitch optimization
    - Photon Detector design optimization
    - Put Anodes on the Outside
    - Signal and Background Requirements for Supernova burst and proton decay physics
  - FD studies postponed in favor of 35-ton work. Personpower permitting, these can be reinvigorated.

Scientific Goals for FY15 and FY16

- Direct effort toward GEANT-based ND simulation
- Refine physics sensitivity calculations
  - Incorporate more detailed systematic uncertainties from cross sections and nuclear modeling
  - Include new detector configuration ideas from the new international collaboration

Cosmogenic Group Goals from Vitaly Kudryavtsev

1. Muon-induced background for proton decay and astrophysical neutrinos. Full simulations, event processing and reconstruction.
2. Neutrino-induced background for proton decay. Full simulations, event processing and reconstruction.
3. Cosmogenic activation as background for low-energy neutrinos.
4. Development of background rejection algorithms.
5. Accurate study of the topology of muon-induced events in LAr.

Large Scale or out of ordinary computing needed to complete these goals: 35-ton

- Need permanent storage of order 100 TB of raw data from the 35-ton prototype. Data-taking choices will affect this number.
- Would like a backup copy of the raw data. Partner labs can host additional copies.
- Need of order 300 TB for storing reconstructed versions. Choices made during reco affect this number.
- May need to iterate reco several times. May overwrite older reco sets if need be.
- Need 2M hours of CPU to reconstruct all data per pass (assumes: 30 Hz trigger x 2 months * 1 min/event reco) Many analyses can be done with (small) subsets of the data.
- 2 GB virtual size is fine for 35-ton simulation and analysis jobs

Large Scale or out of ordinary computing needed to complete these goals: Far Detector

- Far Detector simulations take more than 2 GB of virtual size. 4 GB works.
- Most FD studies use a smaller geometry for convenience
- Interactive testing with 4 GB size suits most of our needs.
- >= 100 slots with 4 GB capability would be appreciated but not required.

Did you meet your FY14 Scientific Goals?

- Yes, sufficient progress was made on our scientific goals, but we are by far not done.
- Most goals on ELBNF are long-term
- Software and Computing DOE review May 2014 stressed and clarified many of our goals and deliverables.
- DOE recommended that we
  - develop a computing plan: schedule with resource requirements
  - establish a recurring review process
  - Evaluate our computing project management and financial structure
  - direct more effort towards 35t
- And then the P5 report came! We continue with our physics program as is but some plans are premature.

Did you meet your FY14 Scientific Goals?

- Beam Simulations
  - Optimizations performed – L. Fields did a great job
  - Optmization used FastMC to connect beam flux to physics sensitivity output.
  - Bug found in FastMC parameterization of yields. Need to re-do some steps of optimization
  - 1.2 MW target and horn designs studied

- Success! But more work needed. Explore target and horn designs beyond the NuMI defaults.

- MARS simulations used to model and improve design of hadron absorber, shielding, and beamline components

Did you meet your FY14 Scientific Goals?

- FD Simulation and Design
  - In flux! New international collaboration brings in new ideas for the FD.
  - We continue to refine the design for the 1-phase modular detector
    - Wire angle study – move to 1 wire wrap with lower angle to reduce ambiguity
    - Cavern shape optimization
    - e-gamma separation with ionization work
    - Cosmic Tracker and stitcher between APA's
    - PANDORA
    - Reliable efficiency, PID and resolution numbers not yet available
  - Most sim/reco effort directed at 35-ton as per DOE recommendations

Did you meet your FY14 Scientific Goals?

- Near Detector – preliminary straw-tube tracker simulation now in place.

LBNE Batch Usage with FIFEMON – May 1, 2014 – January 26, 2015

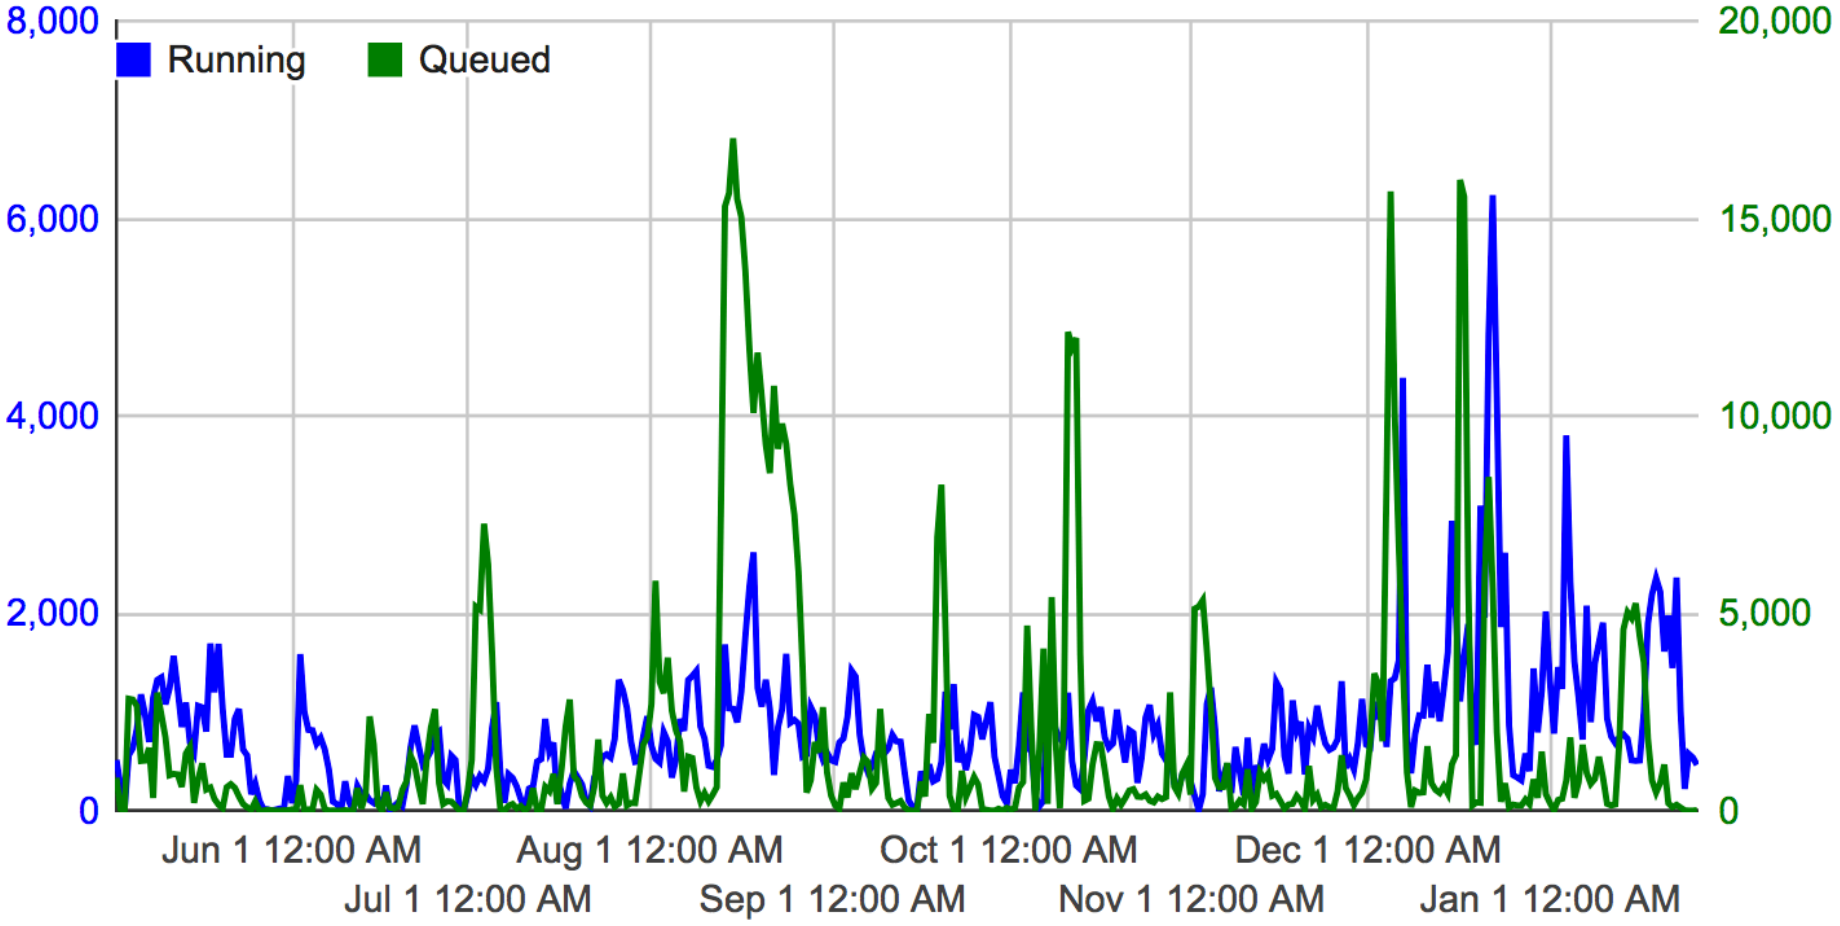

LBNE VO CPU Usage by User from Gratia Monitoring

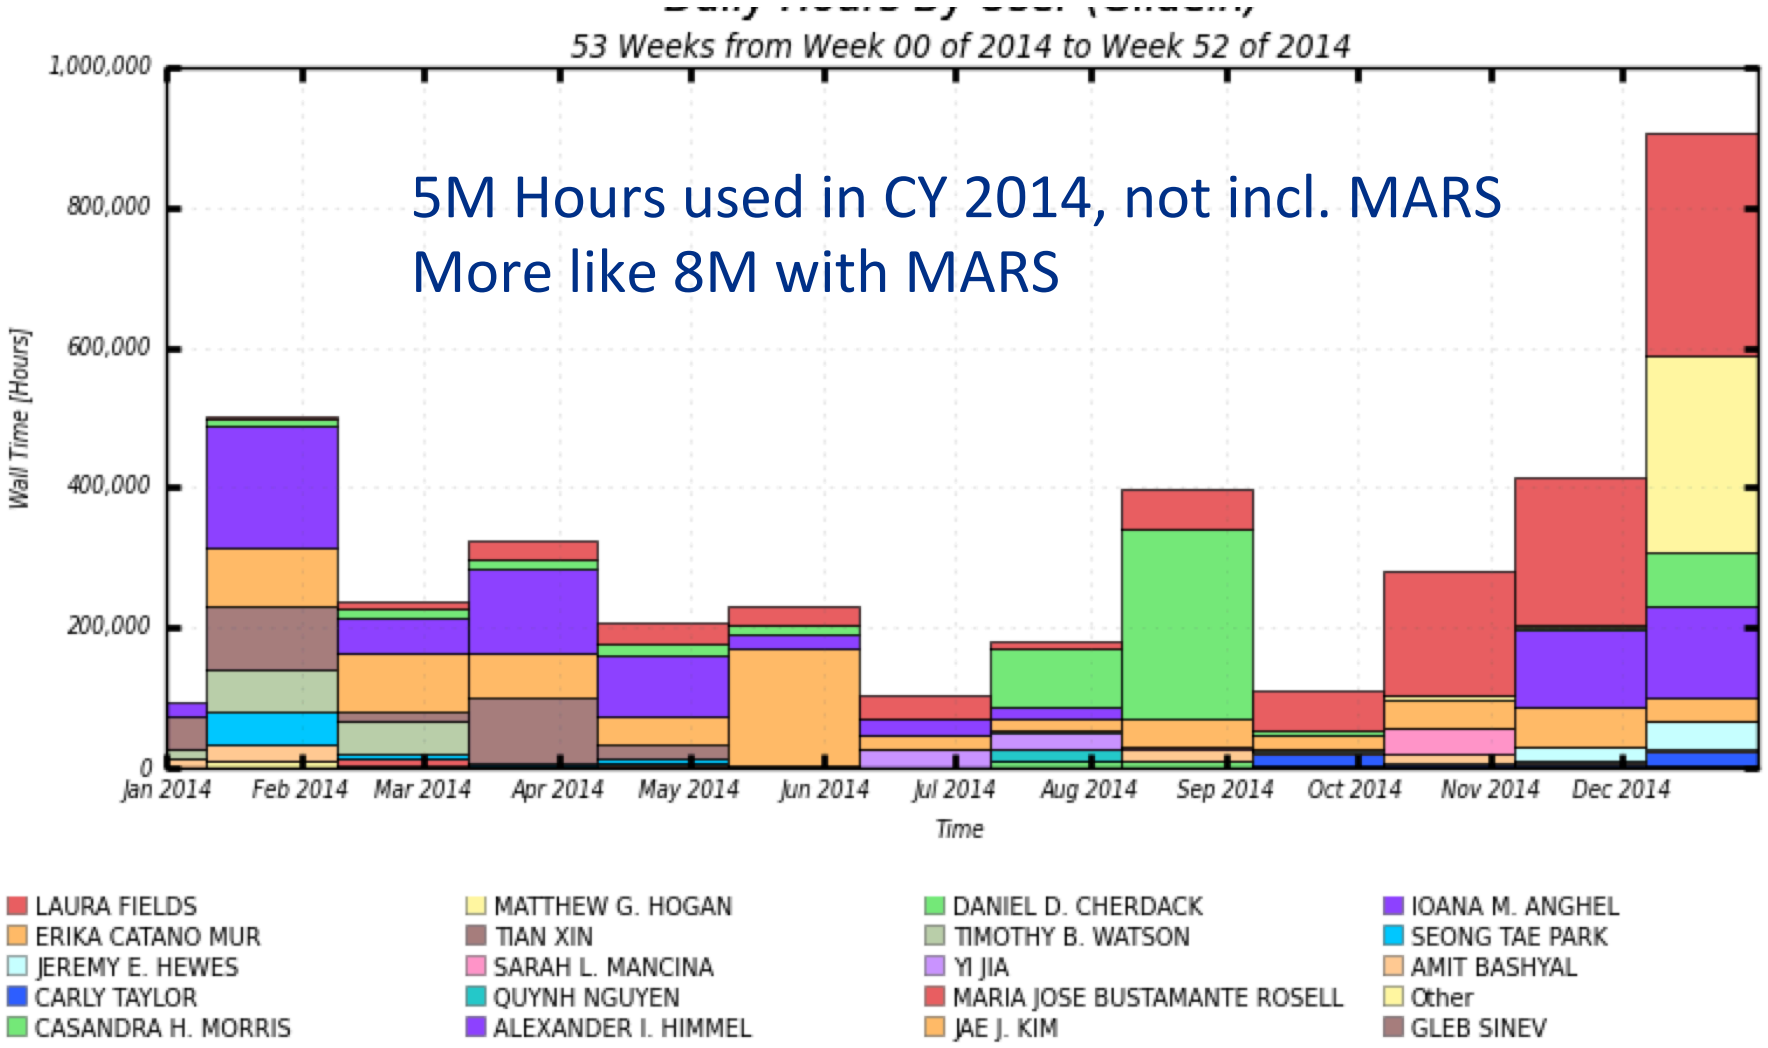

NAS Disk Usage for the Past Year

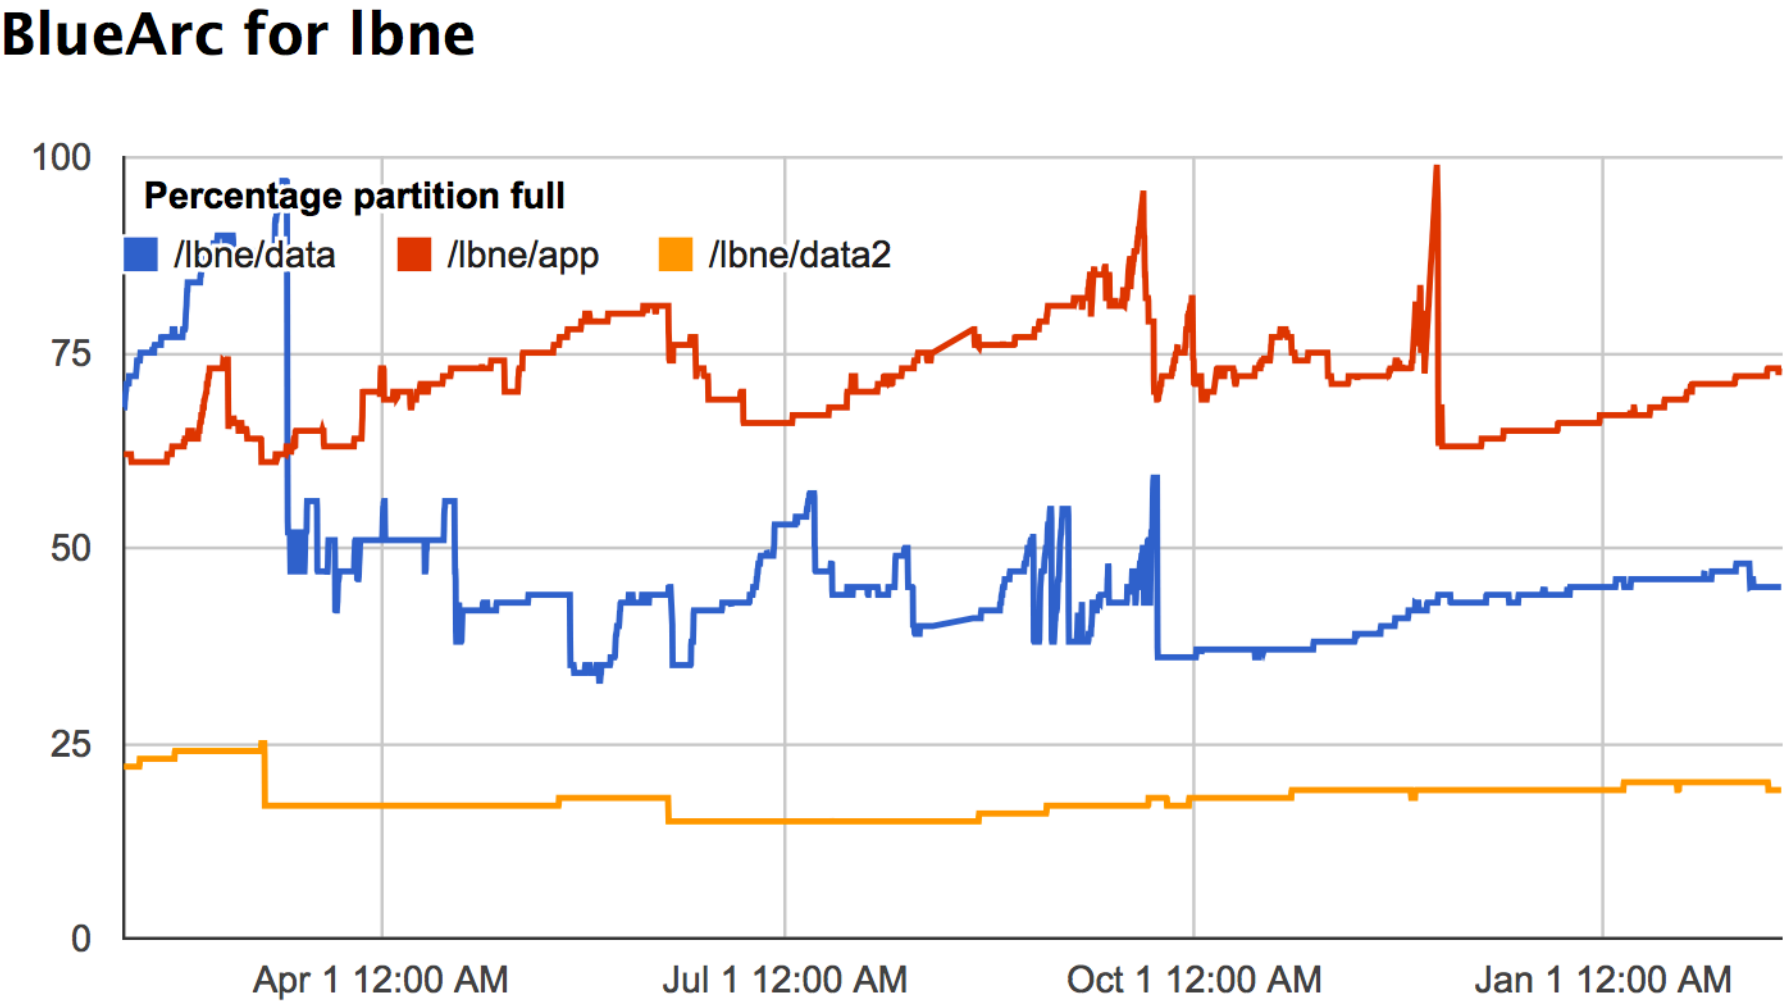

Current Typical Workflows by Group

- Beam simulations and FastMC have similar workflows
  - Run G4LBNE or FastMC on the grid. Store trees on BlueArc and dCache scratch space. Make plots with root.
  - Run standalone programs on the grid using trees (optimization).
  - If storage gets tight, we can put older sets on tape and access via SAM (not happening yet that I know of).
- FD/35-ton sim/reco
  - Run larsoft jobs on the grid – store art-formatted rootfiles on tape using SAM, dCache, BlueArc, and local scratch disk.
  - Analyze art-formatted output files in further larsoft runs using SAM to access data on tape (not heavily used yet).
  - Need adequate staging disk space to analyze data. But we do not expect many iterations through the full reco set, mostly smaller pieces
- ND sim/reco
  - Still a new effort; reliant on BlueArc. Will use other storage as users get up to speed.

Current Typical Workflows by Group

- Cosmogenics
  - Largely a University effort. Their storage requirements are almost all of our BlueArc space so they will have to use tape and SAM if they want long-term persistency at Fermilab
- MARS
  - CPU-dominated jobs, modest storage requirement. BlueArc should suffice

- A note on personnel availability: Many collaborators work on ELBNF part-time. They work on something else for a few months and come back wanting their stuff to be there.

  So far, persistency of data and code hasn't been an issue.

  But cache lifetimes of <2 months would not be enough.

Summary: FY15 Requests: Batch and Storage

Group	CPU Hours	Disk (TBytes) Permanent	Tape (TBytes)	Peak Slots
Beam Simulations	2M	20		2000
MARSLBNE	3.6M	8		1000
35-ton, FD Sim/Reco	4M	10	400	2000
ND Sim/Reco	0.6M	4		500
Cosmogenics	0.5M	10	50	200
FastMC/Physics WG's	3M	7		2000
DAQ	modest	incl in 35t		
Water Cherenkov	1M	0.5		120
Total	14.7M	59.5	450	6000

Storage is not cumulative, just absolute need
Simultaneous peak demand can occur.

Summary: FY16 Requests: Batch and Storage

Group	CPU Hours	Disk (TBytes) Permanent	Tape (Tbytes)	Peak Slots
Beam Simulations	2M	20		2000
MARSLBNE	3.6M	8		1000
35-ton, FD Sim/Reco	4M	10	400	2000
ND Sim/Reco	0.7M	6		500
Cosmogenics	0.5M	10	50	200
FastMC/Physics WG's	3M	7		2000
DAQ	modest	incl. in 35t		
Water Cherenkov	0.1M	0.5		120
Total	13.9M	61.5	450	6000

Storage is not cumulative, just absolute need
Simultaneous peak demand can occur

Needs from Service Areas: Grid and Cloud Computing

- Request approx 15M CPU hours in each year, FY15, FY16
- We would like to be able to use up to 2500 or more slots at a time (peak usage).
- We have the ability to run offsite and have tested jobs that do useful work on the OSG
  - FD/35t and ND sim/reco require maintained releases of larsoft,art,lbnecode,hisoft and depndencies on CVMFS to run offsite
  - FastMC and Beam simulations port easily – tarballs.
- MARS: modest resource allocation at NERSC, applied for ANL BlueGene time. Nikolai is performing tests on ANL's ALCF with T. LeCompte

Offsite Computing Opportunities

- Collaborators frequently advertise availability of their university or group computing resources
- Frequently these are small (<< Fermigrid)
- OSG provides a set of uniform supported features we can expect to be present on worker nodes.
- More resources available for less effort if configured for OSG
- But can be a good project for international partners with some time, people, expertise, and willingness to learn to adapt even a small resource to collaboration use.
- More likely, self-contained projects can be run at university facilities. R&D projects, standalone physics tasks.

Needs from Service Areas, FIFE, and others

- Networked Storage
  - BlueArc: 60 TB data, 2 TB app area are OK. Quotas help
- Interactive servers: Six 4-CPU VM's seems adequate.
- DAQ Test VM
- Database Servers
- GPU-enabled interactive servers (2? For software devel.)
- Grid and OSG computing support, documentation
- Batch job and storage monitoring. FIFEMON/Gratia/Others
  Would like FIFEMON maintenance and improvements.
- VOMS
- Physics and Detector Simulation: GENIE, Geant4 support

Needs from Service Areas, FIFE and others

- Scientific Data Management
  - SAM – Many thanks to a large team!
  - Enstore (tape), dCache (tape-backed and scratch)
  - Most of our 500 TB storage request should be on Enstore and dCache
  - IFDH
  - xrootd
  - POSIX access to dCache very useful
  - Authentication issues when accessing data offsite
    - Would like this to be easier. Special requests made for individual users seemed harder than it should be
    - Kerberos-based grid proxies for most people.

Needs from Service Areas, FIFE and others

- Scientific Frameworks and sofware tools (development, documentation, support)
  - art: many thanks to the art developers!
  - artdaq
  - larsoft: special thanks to larsoft developers!
  - git, svn (cvs for the holdouts)
  - Redmine
  - mrb

Needs from Service Areas, FIFE and others

- Jobsub
- LXR, Doxygen, other source documentation/search tools
- Totalview
  - license, installation on gpvm's
- CVMFS
- Databases, online and offline.
  - Experiment-specific parts being developed by the collaboration and coordinated within the larsoft group to make consistent with other experiments. Larsoft group is integrating efforts from different developers to have a single solution.
- Jenkins/Continuous Integration
- Metadata – using tools from MicroBooNE adapted to LBNE for 35-ton work. Other groups not using tapes yet.

Needs from Service Areas, FIFE and others

- DAQ and Engineering
  - use artdaq
  - Would like some surplus rackmount/desktop CPU's for use at PC4 for the 35-ton prototype test. Need a Fermilab-managed computer to run online DQM jobs
- Document Management
  - DocDB
  - Indico
- Project Management
  - Primavera
- Visualization tools (esp. Beam Simulations)
  - licenses
  - support

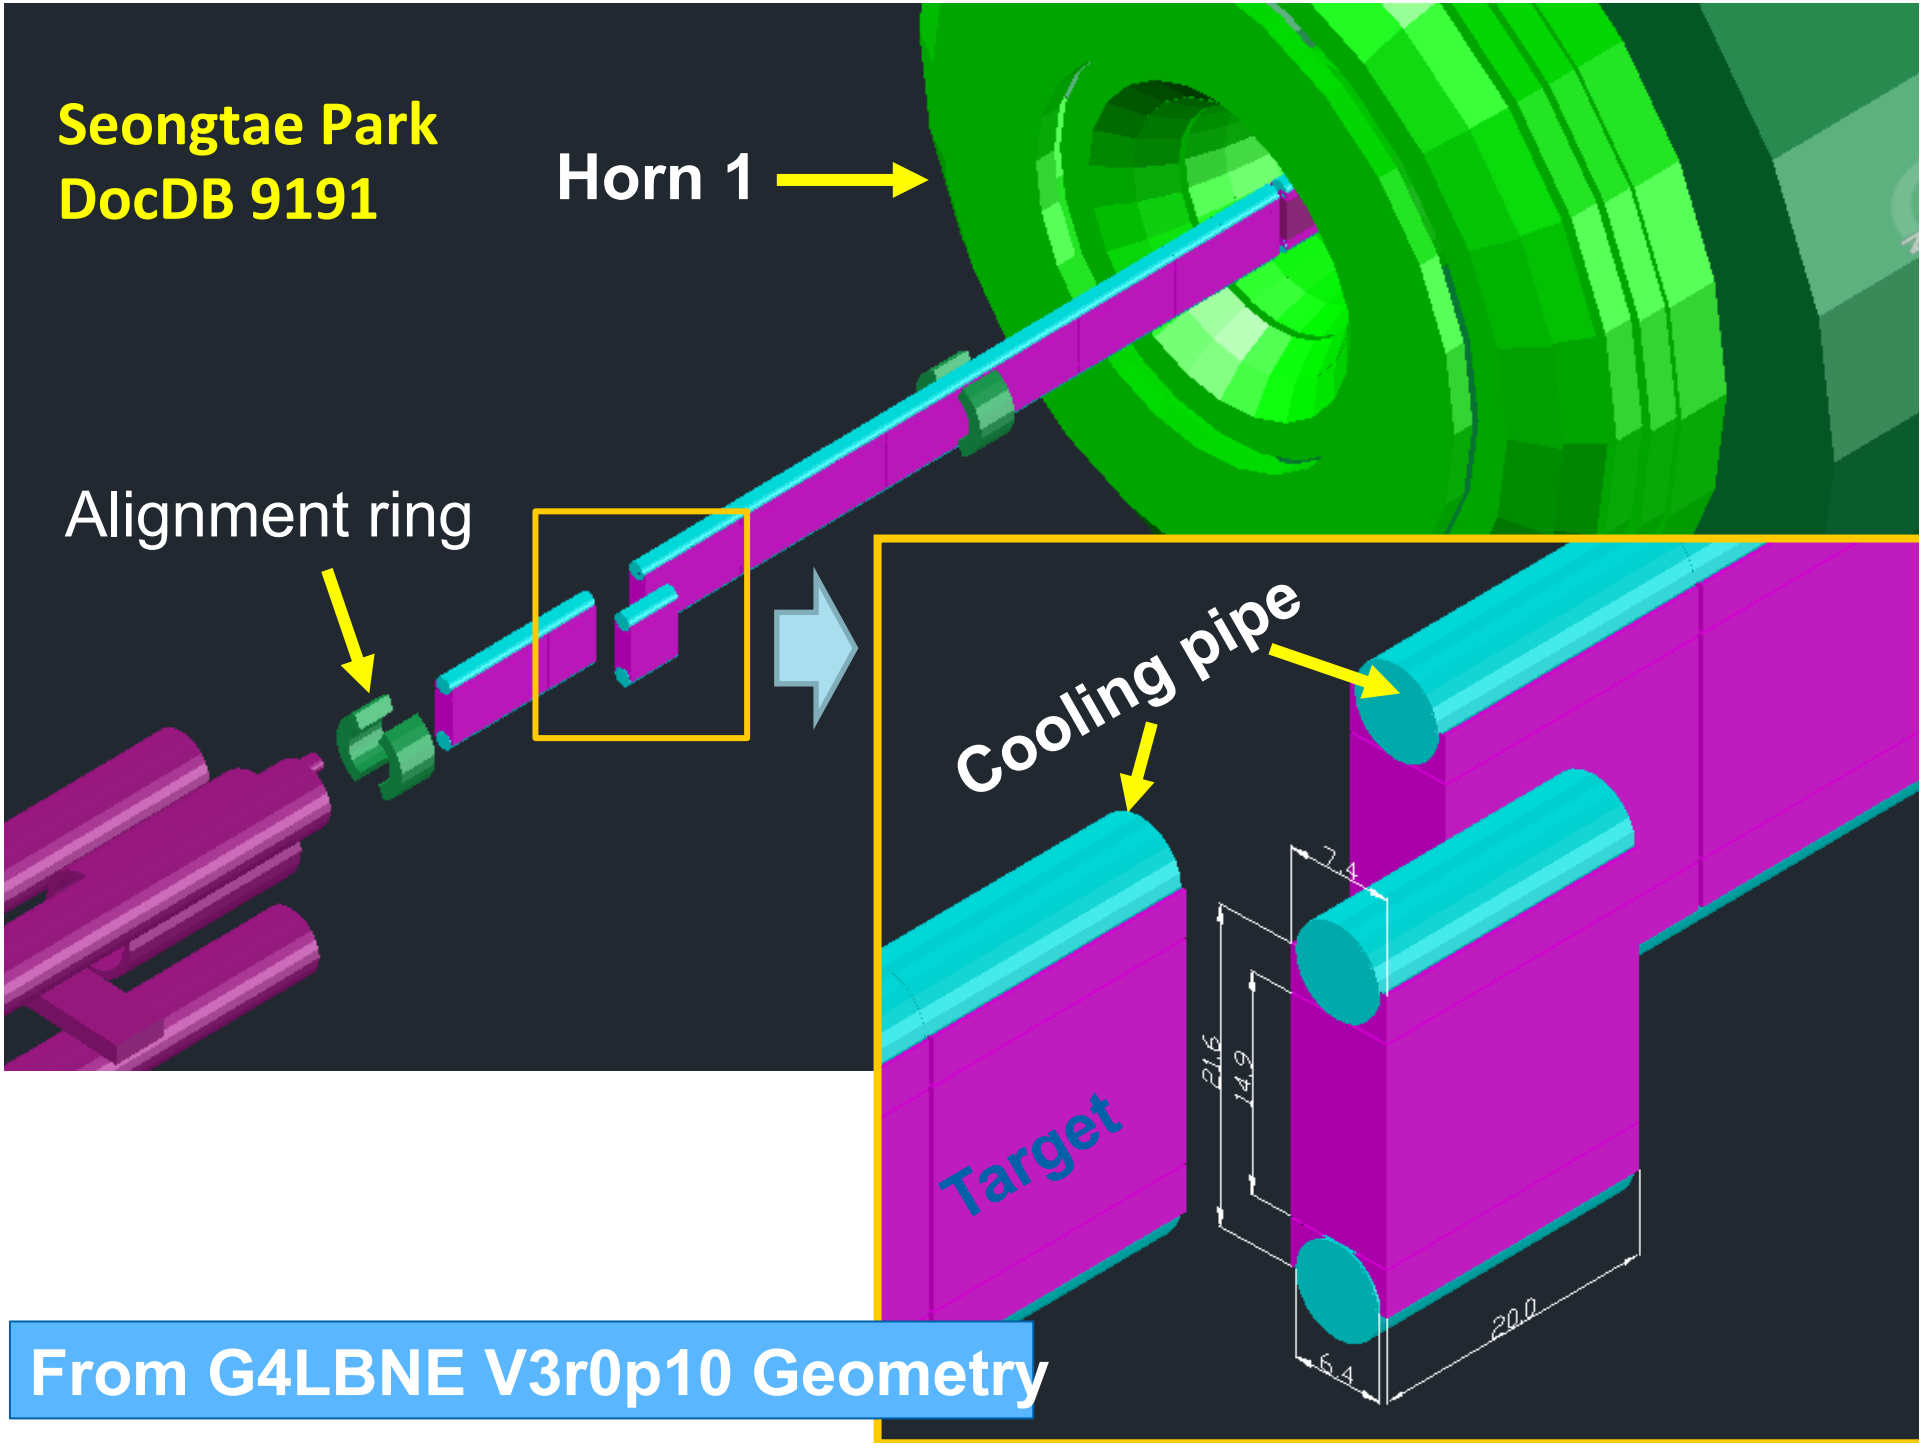
Seongtae Park
DocDB 9191
Horn 1
Alignment ring
Cooling pipe
Target
2.4
21.6
14.9
20.0
6.4
From G4LBNE V3r0p10 Geometry

Needs from Service Areas, FIFE and others

- Central Web Hosting
  - Experiment homepage
  - Sharepoint
- AFS home areas and web areas
- At least one general FNALU node is useful for administrative purposes. (I use flxi02).
- Conferencing tools
  - ReadyTalk works, but has limitations
  - Vidyo would be nicer..

TSW/EOP Status and Plans

- Are your TSWs signed and up to date?
  - Meeting scheduled to start this process
- If not, do they need revision?

Future Directions (Challenges and R&D)

- Changes in SOP
  - New International Collaboration!
  - New Beam and Detector Ideas
  - New Working Groups
  - 35-ton Prototype Assembly, Commissioning, Processing, and Analysis
- Future R&D Projects
  - CERN Test Beam
  - University R&D: Domestic and International
  - CAPTAIN (separate item, not included in this request)
  - LArIAT (separate item, not included in this request)

Batch Usage – Last Two Weeks according to FIFEMON