Development of the LBNE Computing Model

Liz Sexton-Kennedy

Maxim Potekhin

LBNE Software and Computing Coordinators

Overview

We will present and discuss the following topics:

- Current structure of the LBNE Computing Effort
- General direction of its evolution in 2013-2014
- LBNE Software and Computing Requirements – the document
- What we learned in DOE review of LBNE Computing in May 2014
- What we learned from the P5 report
- What all of this means for the LBNE Computing Model
- LBNE and FIFE
- Our next steps

LBNE Computing: Org Structure and Requirements

- Organizational structure: there are the "Software and Computing Organization" and the "Physics Tools Group" respectively – S&C and PT for short, which work in coordination with the DAQ/Online Group (E. Berman on the Project side, others).
- Tom Junk is the leader of the Physics Tools Group and myself and Liz Sexton-Kennedy are heading the S&C.
- This presentation will mainly cover the S&C (more detail in a couple of slides).
- The Computing Model: its main function is to fulfill the requirements and needs of LBNE. It is important to for a consensus regarding what's needed.
- For these reasons, we have developed the Software and Computing Requirements (DocDB 8546) - based on which the LBNE Computing Model shall be created. The Computing Model will serve as a reference for policy and technology choice decisions.

Mission Statement of the LBNE Software and Computing Organization

The goal of the LBNE Software and Computing Organization is to make all aspects of computing technology, software tools and resources easily accessible to the largest number of Collaboration members possible, in a most efficient manner.

LBNE S&C: what does it need to provide?

- S&C aims to provide the Physics Tools Group and Collaboration:
  - All aspects of software management (architectural compliance, revision control, validation, continuous integration…)
  - Database design and operation
  - Data handling system (storage and distribution…)
  - Distributed Computing capability (distribution of binaries, payload distribution to Grids and Clouds, monitoring…)
  - Geometry management and services
  - Visualization tools
  - Networks (performance monitoring, feedback into data and distributed computing systems, near site to far site communications etc)
  - Collaborative Tools (documentation databases, content management, conferencing etc)
  - Cybersecurity policies and procedures (in cooperation with site personnel)

Status and Progress in 2013-2014 (1)

- LBNE computing has been centered at FNAL since 2012. FNAL has provided most computing resources. According to T.Junk, non-FNAL sites contributed computing resources on a smaller scale (but no federation of resources).
- Substantial levels of support at FNAL allowed LBNE members to be productive early on with low upfront costs. Significant investment in LArSoft (ArgoNeuT, MicroBooNE) was leveraged for the benefit of LBNE, and some challenges were inherent in that (code refactoring etc).
- The is a strong coupling of LBNE software infrastructure to that of FNAL, which in some cases resulted in lack of flexibility in software configuration and accessibility. Building LBNE software outside of FNAL was difficult, effectively presenting a barrier to wider engagement of LBNE scientists in Physics Tools development and other aspects of R&D (close to resolution in 2014, see next slide).
- Expertise and R&D contributions were largely distributed while the computing resources were largely centralized.
- In early 2013 revision control was done mostly with SVN (but migrated to git by 2014).

Status and Progress in 2013-2014 (2)

- Useful interaction with the *art* team at FNAL (extensive communication and meetings in person).
- Explored technologies and approaches for "build orchestration" necessary to implement a portable and robust set of procedures for building and installing LBNE software on supported platforms anywhere.
- Identified the "*waf*" build automation tool as a promising candidate for the core build engine – used by many high-profile projects in industry in research.
- Created a configuration management layer on top of *waf* to simplify and systematize its usage (codename "*worch*" for *waf* orchestration). Can be used to build and configure virtually any software or application. More in a talk by B.Viren.
- Distributed Computing – see next slide.

Distributed Computing

- Background:
  - The FIFE project at FNAL started providing facilities for job submission to the Grid outside of Fermilab in 2013. Then, usage by LBNE remained low for various reasons (varying levels of resource requirements, space and allocation etc, difficulties in software build and distribution).
  - There is a strong synergy between FNAL, various HEP and other project on one hand, and the Open Science Grid Consortium (OSG) on the other hand. The OSG presents ready tools and systems for resource federation.
  - Valuable work was done by both FNAL and OSG to create and maintain the CVMFS infrastructure (in different enclaves) allowing transparent distribution of software and configuration data to worker nodes over the network, which really opens the door to agile and truly distributed placement of computational workload on the Grid.
- Status:
  - We engaged the Open Science Grid Consortium and utilized our prior experience with the OSG environment (M.Potekhin).
  - A few concrete types of LBNE payload have been identified.
  - CVMFS area has been configured (Thanks to M.Kirby!).
  - Initial test runs placed LBNE payload jobs on 12 sites managed by the Open Science Grid.
  - Exceptional level of engagement and support by the OSG.
- Configuration and Monitoring – TBD.

Data Handling (2013-2014)

- Background:
  - Satisfying the 35t Prototype data requirements is a priority in 2014.
  - Data formats and protocols will be coordinated with the DAQ/Online Group.
  - The Physics Tools Group has identified its initial needs for Metadata.
- Status:
  - FNAL has remained the basis of the LBNE storage capability.
  - SAM is the system in which FNAL invested heavily in the past few years and it will be used to meet immediate needs of LBNE for metadata and catalog (e.g. to satisfy the requirements of the 35t prototype, MC production etc).
  - Plans are being made to have a xrootd service running at FNAL to enhance our distributed computing capability.
  - Tentative plans are being made for placing data replicas at BNL and LBNL, using bulk data movement tools where necessary.

Geometry and Visualization

- Background:
  - As with any advanced particle detector project, geometry modeling and description is important. The Liquid Argon TPC is by its nature a “visual” tracking detector and this capability can be greatly leveraged by introduction of versatile visualization tools.
  - The are a few (mostly application-specific) solutions across LBNE but no unified geometry management system (no central version control).
  - Due to vigorous R&D program, this is more important now than it will be in 5 years down the road.
- Status:
  - There is some pre-existing experience with geometry management and visualization in LBNE coming from both IF and EF.
  - There is progress with event display for the 35t detector and beyond (BNL, Otterbein College).
  - We plan to assess our options for the geometry description and model and implement a solution that can be utilized in simulation, reconstruction and analysis, and visualization systems.
  - We don’t plan to implement any new system for the 35t prototype.

DOE Review in May 2014

- Main comments and recommendations (relevant to S&C)
  - Software and Computing in LBNE will remain off-project.
  - In general, DOE wants us to follow practices of the LHC experiments (ATLAS and CMS) which are considered a success.
  - LBNE needs to develop a Computing Model and provide sufficiently developed metrics (e.g. quantitative measures to guide the planning process). We were already planning to complete the Computing Model within a year.
  - In general, DOE wants to have “detailed plans” (hard to do with no guideline for resources to be provided!). Need a prioritized list of tasks, dependencies, deadlines.
  - A review process will be established to update computing requirements etc.
  - Concern about engagement (i.e. how new collaborators can join the software effort in a productive manner). We need a plan for that.
  - Fractional FTE involvement considered not optimal.
- Bottom line – we need to create the Computing Model, backed up with quantitative measures, and develop a compelling plan of work with priorities and a timeline.

P5

- The scope and scale of LBNE are to increase.
- There will be an international *facility* for IF experiments.
- For software and computing, we need to assume that for the next months it's business as usual – we are still under obligation to deliver items per the review recommendations, and also need to improve the infrastructure to help in current PT work – and to facilitate engagement. This is in line with the project becoming "more international".
- Same follows from the timeline – next review will happen in a few months, and any significant evolution of LBNE into LBNF is likely to take more than that.

Technology trends

- In the past decade we witnessed a rapid evolution of computing in both industry and science domains
  - Cloud computing matured and became pervasive (and useful)
  - New database technologies offer unprecedented scale of data storage and manipulation
  - Continued trend to implement a fully distributed model for virtually every type of computing resource, from the network itself (cf. SDN) to database (distributed noSQL) to CPU (Cloud)
  - Build automation and Continuous Integration became a staple of virtually every software shop
- We have valuable experience of large scale EF and IF collaborations to learn from.
  - “Flattening” of the computing model in HEP experiments, i.e. transition from a strictly tiered approach to a more symmetrical ensemble of sites
  - Maturing of fully network based methods of access to, and distribution of both software and data, cf. CVMFS and xrootd
  - Demonstrated value of robust Grid monitoring capabilities and Workload Management Systems

Computing Model - examples

- According to the DOE review, Computing Models of ATLAS and CMS could be used as guidance for LBNE. I believe this is in large part due to the ability of these experiments to leverage widely distributed resources, but that's just one factor.
- Of course there are plenty of other successful projects (e.g. in the Intensity Frontier domain) which we can also learn from, such as Daya Bay and others which achieved a remarkably short time-to-results.
- Broadly speaking, there are common characteristics we observe in these computing models. They also appear to satisfy the LBNE Computing Requirements.

Computing Model: desired characteristics (1)

- Software infrastructure: rigorous version control, build system, release management, validation and documentation. Continuous integration (CI). Specific details of this process may vary (cf. more rigorous code review in CMS compared to ATLAS), but the principles are similar.
  - Recent example: we currently need to invest effort in LBNE code management for LArSoft-based software and ensure compatibility with the evolving LArSoft suite. David Adams (BNL) recently joined this effort. A CI system would have saved us a significant amount of work
- Software portability and installation: it is possible to build the complete suite from source anywhere. There must be automated mechanisms of pre-built software delivery to sites.
- Distributed data: fairly symmetric data placement strategy which doesn't have a single "go to" location for most types of processing to take place.
- Distributed CPU: while computing centers vary in node count, core count, memory size and other characteristics, there is no single data center (outside of the primary raw data sink) that is 100% critical for data processing to take place.

Computing Model: desired characteristics (2)

- Workload management:
  - In most cases, there is at least an “insulation layer” on top of the Grid infrastructure to make managing workflows more efficient.
  - Jobs can be submitted from any location to any site belonging to the federated resource.
  - Monitoring capabilities in such systems are very helpful and boost productivity.
- Recent advances – already incorporated into evolving computing models:
  - CVMFS (code in the Grid)
  - XrootD (data in the Grid)

LBNE Computing Model and FIFE

- We believe that the LBNE Computing Requirements (DocDB 8546) do reflect what the experiment needs in the area of software and computing.
- The value of the FIFE toolkit(s) for LBNE is determined by that.
- As noted on previous slides, the LBNE computing model will include
  - Continuous integration infrastructure
  - Distributed data (multiple storage elements, “data in the Grid”)
  - Distributed processing (multiple computing centers) and agile workload management, user-centric monitoring (as opposed to “site status” type of monitoring)
    - Ability to direct workload to any of the multiple federated sites as defined by resource availability and allocation. Example – researcher at LBNL must be able to submit jobs to University of Colorado, BNL etc, transparently.
    - Ability to assess the status of tasks and jobs at varying degree of granularity, down to a single job, with complete error logging and diagnostics
- We would like to understand whether FIFE can provide these capabilities.