

LBNE Software and Computing Requirements (draft 1.1)

LBNE Software and Computing Organization

February 1, 2014

Abstract

This document presents the high-level requirements that the LBNE Collaboration has pertaining to its Software and Computing (S&C) effort. These Requirements are intended to state problems needing solutions rather than describe any particular solution itself. Specific implementations or technology choices are explicitly not included here but will be provided in future documentation.

1 Introduction

1.1 Purpose

Analyzing and addressing these Software and Computing Requirements will result in the creation of the LBNE Computing Model. The Computing Model document will then be used by the Collaboration to steer its S&C Organization and to evolve its computing platform in the direction optimal for achieving the scientific goals of LBNE. The Computing Model and documents pertaining to its implementation will also serve as instruments to inform the funding agencies about the scope of work and rate of progress of the LBNE Software and Computing Organization.

1.2 Origin and Scope

These Requirements are developed based on "Physics Tools and Software and Computing Organization of the Long Baseline Neutrino Experiment: the goals, the structure and the plan", FNAL DocDB 7818. They serve as a step towards the LBNE Computing Model. Upon acceptance by the Collaboration, the Computing Model will be considered a normative basis for policies to be implemented by the LBNE S&C Organization.

This document does not cover functional requirements for the Physics Tool software which is being developed by the corresponding organization in LBNE (one example of such functional requirements would be the characteristics of the tracking system or physics performance of the Monte Carlo simulations). In accordance with the Plan mentioned above, the S&C Organization focuses its effort on the foundations of the Software and Computing infrastructure, such as code management, Grid capability, data management etc (and does not concern itself specifically with the application code such as tracking, analysis etc).

An effort was made to maintain a relatively high-level view of the S&C issues, and to not go into smaller details which are more likely to change as the project moves forward, while still providing adequate basis for making informed decisions. In cases where it was impossible to establish concrete metrics or parameters for a specific requirement, it is still listed as an item the S&C Organization will need to address in future.

Information and opinions necessary for the creation of these Requirements have been collected during a number of LBNE Software and Computing meetings, conference calls and extensive information exchange via e-mail and other means. It is recognized that the Requirements themselves (and the Computing Model) may evolve with time, in order to correctly reflect the status of rapidly evolving technologies and the LBNE organization itself. Such expectation is reinforced by actual experience of large scale HEP collaborations. We anticipate therefore that the requirements will be revised roughly in the middle of the time period between their creation and the commissioning of the experiment, and help evolve the LBNE Computing model such that it continues to meet the needs of the Collaboration.

1.3 Document Structure

The document is organized as a set of categories reflecting major responsibilities of the LBNE S&C Organization. Each category will generally contain information including:

- **Description:** Description and scope of the category.
- **Definitions:** If necessary, specific definitions of certain items are given.
- **Issues:** Statement of issues that requirements in this particular category are expected to address.
- **Requirements:** The requirements themselves, formatted as assertions.

Where necessary, each of these items may be placed in an individual subcategory to better reflect details and granularity of information being presented. A requirement section may be accompanied by one or more of the additional items:

- **Use cases:** Descriptions of any use cases that drive the choice for the requirement.
- **Justifications:** Explanations as to why the requirement was chosen.
- **Accepted risks:** Any known problems with the requirement which are considered acceptable.
- **Alternatives:** Any alternatives that were considered with reasons for rejection.
- **Deviations:** Recognized cases where deviations do not violate the intention.

In addition, in cases when it helps clarify the scope of a given requirement, appropriate statements will be made as to what is **not** required in the context of a particular section.

2 Software and Computing Organization of LBNE

2.1 Description

This section contains a few requirements related to the structure and function of the LBNE S&C Organization itself. While not covering the S&C organizational structure as a whole, the items in this section need to be clarified in order to provide a basis for certain statements in the rest of the document. This section is different from others in that its suggestions are more concrete as they relate to the S&C organizational structure.

2.2 Issues

- Need to identify and govern technology choices and system architectures in an environment with widely distributed computing facilities of varying capability and characteristics, maintained and utilized by personnel in geographically diverse locations.
- Need to find balance between small scale (or local) convenience and large scale (global) functionality.
- Necessity to have a definitive source of technology policy and decision-making in the Collaboration.
- The complex distributed computing infrastructure to be utilized by LBNE will eventually require a focused effort of a group of experts, providing ongoing operational support.

2.3 Technical Advisory Committee and forming Consensus

- a Software and Computing Technical Advisory Committee (TAC) shall be formed. As its name suggests, its role is advisory and not normative.
- initially, the TAC will consist of five members selected by the spokespeople. Coordinators of the S&C Organization, the Physics Tools Group and the Physics Working Groups Organization will be additionally included as members ex-officio.
- interfaces with software and computing elements managed by the LBNE Project shall be considered in the Committee recommendations.
- the purpose of TAC is to advise the LBNE collaboration spokespeople on technical matters regarding collaboration software and computing. The committee is to proactively identify, examine and select technologies pertinent to LBNE software and computing. Recommendations by the TAC will be reached through committee consensus.
- important decisions will be made by consensus. Consensus is to be formed by discussions within the S&C Organization and in consultations with TAC. If this can not be achieved then the TAC will be asked to provide its opinion to LBNE spokespeople, who will then make decisions based on the information provided to them.

2.4 Operations

- The Software and Computing Organization shall form and manage the Operations group, whose mandate will be to support ongoing LBNE data processing, including but not limited to data management system, database performance and availability, status of LBNE Grid sites etc.
- LBNE Operations will work in close cooperation with site personnel.

2.5 Justifications

- Due to considerable size of the Collaboration, it is important to identify a group of individuals who possess sufficient expertise and can also firmly commit to creating recommendations and conducting evaluations in the Software and Computing sector of LBNE, while doing it in an expedient manner.

- The ultimate decision-making power rests with the spokespeople of the Collaboration, which under normal circumstances also necessitates formation of expert opinion on technical matters.
- Smooth functioning of LBNE computing "plant" will depend on agile support by experts, whose availability for this task must be guaranteed at any given time, hence the Operations group.

2.6 Alternatives

If the TAC is not formed, the S&C Coordinators may delegate certain responsibilities to expert committees as necessary, to address specific issues of the LBNE S&C effort.

3 Data

3.1 Description

This section contains the requirements related to handling a variety of data by LBNE Collaboration. This includes data produced by LBNE primary detectors and monitoring systems, as well as derived (processed) datasets and other collections of data. Databases and related technologies will be covered separately in Section 4.

The principal requirement for data distribution and access methods is that they need to support a coordinated and widely (indeed globally) distributed network of computing centers and research groups located at LBNE member institutions. This is closely related to Section 6 (Distributed Computing).

Data Retention policies will aim to provide cost-effective and optimal schedule of data distribution, replication and retirement, in order to maximize efficiencies in achieving the scientific objectives of LBNE. On a longer time scale, there is a need to put in place policies and procedures for data preservation. The principal difference between retention and preservation is that the former concerns itself with optimizing ongoing analyses and other types of data processing, whereas the latter addresses the long term storage and documenting and preserving methods of accessing these data (formats, algorithms, application code etc).

3.2 Definitions

- **Raw Data:** Data saved by a detector (far, near or prototype detector) DAQ or monitoring systems and information from FNAL beam monitoring.
- **Processed Data:** any data produced by production software for the collaboration as a whole or to satisfy the needs of working groups. This includes simulations, data reduction/reconstruction, data skimming and inputs to final analyses. This type of data does not include data samples produced by individual users using software that has not been certified as production software by appropriate Working Groups, Production Managers and if needed by the S&C Coordinators.
- **Production Software:** the suite of software run to produce an official collaboration result. One example of such data is a data sample used in a publication. Such software is subject to strict version control, QA and validation, and is utilized in a managed fashion.
- **File catalog:** a general term used to describe a system which performs a range of mapping functions, such as mapping of a Logical File Name (LFN) to one or more physical locations of the file in the distributed data management systems. This functionality is essential for locating physical copies of the data when needed, optimal matching of distributed data to available computing resources, storage accounting and various other aspects of distributed data management.

- **Metadata:** data that describes other data.
- **Dataset:** a collection of files (potentially in differing formats) and corresponding metadata (uniform across the set) that forms a coherent unit of data used in a computation, and is accounted for as such.

3.3 Issues

The following is a short list of issues that LBNE S&C Organization will need to address in its data handling sector:

- Data replication strategy for each major class of data needs to be created.
- Data retention policies are an important tool to assure cost-effectiveness and overall efficiency of the S&C data infrastructure.
- Long-term data preservation policies and procedures need to be properly planned.
- General rules of access for LBNE collaborators, and for general public in order to assure efficiencies and compliance.
- Data distribution and access methods will play a major role in resource availability and overall efficiency of resource utilization.
- The file catalog is one of the central elements of many architectures for handling data, and its performance characteristics are of utmost importance.
- File and dataset metadata: management of the large and heterogeneous volume of data in LBNE (e.g. Monte Carlo samples, raw data, processed data) in any stage of analysis or transformation etc) requires creation, storage and appropriate use of coherent metadata, which must allow identification, location and retrieval of collections of data necessary for specific purposes. It is crucial that the design of the metadata and any system for its utilization is such that it allows for truly distributed, highly scalable and symmetric strategy of data placement.
- Data design and formats: having a consistent approach to data design and policies to maintain relevant standards and interfaces is crucial for efficient and reliable software development process and operations of LBNE.

3.4 Replication Requirements

3.4.1 Raw data replication

- At any given time, there will be a complete set of all raw data stored at least one principal facility possessing appropriate storage capacity, network connectivity and bandwidth as well as other relevant infrastructure characteristics.
- There will be two additional replicas of raw data kept at locations separate from the principal data storage, for a total of 3 copies of raw data (see "Justifications").
- Each additional replica will not be required to reside in a single storage facility and may be divided in managed datasets distributed to a few designated LBNE computing centers chosen based on agreements with member institution and taking into account their infrastructure characteristics.

- Latency of replication of raw data from the primary copy to the remote sites will be required to be within 24 hour period, counting from arrival of new data to the primary storage location, and to the point where transmitted data is deposited in mass storage at the remote location, validated and accounted for in the data handling system.
- Any storage system or host selected for the official copy (replica) of the raw data will employ storage technology with an expected loss of no more one unit of data per million per year.

3.4.2 Processed data replication

- the data will be generated at, and distributed to LBNE computing facilities at participating institutions based on combination of factors such as research interests of the corresponding working group operating at a particular location, resource availability and scheduling policies of the Workload Management System employed by LBNE
- there will be no fixed minimum number of replicas for the processed data, and the number and placement of copies will be determined dynamically based on requests of a particular processing campaign, with priorities set by a collective decisions of the Physics Working Groups, Physics Tools Group and the Data Management sector of the S&C Organization

3.4.3 General replication

- mechanisms will be put in place to assert validity of the data being replicated and/or transmitted, including but not limited to checksum controls
- an accounting and monitoring system will be put in place (in conjunction with more general Data Handling systems) in order to provide full control over data placement, volume, status and other characteristics

3.5 Data Retention and Preservation Requirements

3.5.1 Raw data retention

- Since these data are irreproducible and irreplaceable, such data will be required to be stored at least for the duration of the experiment, unless a specific and exceptional decision is made by the Collaboration, while storage beyond the time period when LBNE is active falls into "data preservation" category (see below).
- Raw data shall always be readable by software available to the Collaboration, for the duration of its existence, and migration/compatibility tools shall be provided as needed.

3.5.2 Processed data retention

- Specific policies shall be created by the S&C Coordinators and the Data Management Coordinator reporting to them, who will be tasked with collecting information regarding requests for specific data types and segments, monitoring capacity, access patterns and other relevant data for effective decision-making. There will be no fixed retention time set for processed data.

3.5.3 Data preservation

- The S&C Organization will develop long-term data preservation strategy in compliance with regulations put in place by the funding agencies and utilizing best practices in science and industry
- Funding and support model for long-term LBNE data preservation will be established in consultation with the funding agencies by the time of commissioning of the detector

3.6 General Rules of Access to LBNE Data

Access by members:

- Access to raw or processed data by official members of LBNE Collaboration will not be limited by any specific policies.
- Technical implementations must not unduly restrict access by any member of the Collaboration.
- Each member institution and individual member of LBNE is required to abide by the data access and distribution rules contained in this document.

Access to raw or processed data by general public is allowed, subject to the following:

- Unless sufficient external funding is provided, there will be no purpose-built systems to facilitate delivery of most LBNE data to the general public, thus requests will be handled on best effort basis, as the staffing situation and system availability and capacity permit.
- Data being made available to the public will have the same standing as data used in publications, and therefore be subject to similar procedures such as internal review and QA.
- Data shall be made available subject to a disclaimer and liability waiver, to be formulated at a later date.

3.7 Data Distribution and Access Methods

3.7.1 Raw Data

The raw data volume and characteristics make it very distinct from other data types, and it is helpful to consider it separately from the processed data in the context of this section.

- Raw data will be distributed to institutions in a managed manner, based on specific requests, in cases not already covered by the raw data replication strategy.
- Due to sheer volume of raw data, there will be no guarantee of immediate (low latency) access to raw data held in official replicas, outside of actively managed production processing streams.

3.7.2 Processed Data

- There will be a highly symmetrical placement strategy for the processed data, i.e. in principle both input and output data for any job or application can reside at any site or host which is a part of the LBNE distributed data network.
- Accordingly, there will be no fixed requirement to always copy all or any part of the processed data to a single principal location or rely on it as a source of input data.

- Processed data at any location, regardless of its provenance, will be accessible to LBNE users via submitted jobs or applications from any other LBNE site.

3.7.3 All Data

- Direct access to high capacity, high latency back-end storage e.g. tape, from running jobs or applications, will be not allowed.
- When accessed remotely, data will be made available using standard interfaces and protocols compatible with Grid and Cloud technology as implemented in widely available software stacks (cf. the Open Science Grid)
- Provisions will be made to allow utilization of worker nodes or other processing units in cases where extensive local storage is not available, e.g. sites without the Storage Element capability (thus requiring an appropriate data retrieval and upload capability via the network).
- In deciding details of data placement strategy, the S&C Organization will take into account the infrastructure characteristics and support level at participating sites, as well as potential legal or policy barriers impacting systems and individuals based on nationality or location.

3.8 File Catalog Requirements

- A file catalog system will be put in place by the LBNE S&C Organization.
- The file catalog will have interfaces which are flexible and extensible enough to cover the range of data storage and distribution technologies employed in LBNE.
- The file catalog system will cover the totality of distributed storage used by LBNE, i.e. it won't just describe data at a single site, but instead will allow its clients to locate potentially multiple replicas of the data at multiple sites, thus opening optimization possibilities.
- The file catalog will need to have verifiable scalability properties which should allow it to scale up to the necessary level of data volume and throughput, in order to meet the needs of LBNE data processing throughout the lifetime of the experiment.
- At a minimum, the file catalog will incorporate handling of checksum information in order to facilitate detection of errors and sanity checks during operations on data.

3.9 File and Dataset Metadata Requirements

- A metadata system will be created to support the distributed data processing capabilities of LBNE.
- The metadata will cover the data managed by all participating sites, utilizing a variety of middleware and storage media.
- The metadata system will not be coupled to a particular storage solution or architecture, so as to allow flexibility and ways to evolve LBNE data storage.
- Guaranteed scalability of LBNE metadata system will need to be assured according to parameters set by S&C Organization.

3.10 Data Design and Format Requirements

- Both raw and processed data will be created according to formats which have been explicitly accepted for use by the Collaboration.
- Data formats will be developed according to specific mandates set by the S&C Organization and subject to peer review, before their acceptance.
- Simplicity of design and straightforward unpacking and access algorithms will be part of requirements for any data format.
- The S&C Organization will be proactive in evaluation of existing and potential future data formats, and will assure that obsolete and/or inefficient formats are phased out on a schedule that is not disruptive to crucial data processing streams.
- There will be methods put in place for identification of the data format and its version, for any collection of data (e.g. contained in a file).
- Full backward compatibility is guaranteed at any time, for raw data, either via maintenance of legacy interfaces, or by bulk conversion to a new format, with subsequent rigorous validation.

3.11 Use cases

- A working group wants to reprocess all raw data in order to perform searches for new physics or rare events that would be missed by pre-existing standard analyses. To accommodate technology changes over time, the S&C Organization has rewritten very old data into new formats and with each software release has performed validation checks to assure that any changes to the raw data file formats have been accommodated
- A research group from a member institution located outside of the United States requested a large portion of raw data to be placed at their facility, in order to run various preliminary analyses. Upon working out the schedule of data movement, the S&C Data Management sector monitors progress of data transmission which is performed using one of Grid protocols.
- LBNE secured access to a computing facility which unfortunately does not have significant local storage capability. Utilizing "data in the Grid" technology, the Workload Management System makes it possible for the worker nodes at this facility to download input data and upload results to a different location, utilizing high bandwidth networks
- A user needs to locate datasets produced in a few Monte Carlo simulation runs, produced under specific conditions, for further analysis. Finding this data is made possible by the Metadata and File Catalog Systems.
- After a long processing campaign, results of a particular simulation run have been fully analyzed and the campaign declared a success. To conserve storage space, a decision is made to retire the data since the new simulation wave will use an improved version of geometry, which will be adopted going forward.
- A user received a request for a subset of LBNE data from an institution that is not a member of LBNE. According to the rules, the user then refers the requestor to the S&C and general management of LBNE.

3.12 Justifications

- Having geographically distributed multiple replicas of "precious" data is a common and standard way to ensure it's not lost during the normal course of operations where occasional localized technical faults cannot be excluded, and also in case of natural disasters where a single data center may suffer damage and data loss.
- Replication factor of 3 for the primary raw data was chosen by looking at experience of large HEP experiments such as ATLAS.
- While giving access to LBNE data to the general public is an accepted policy that must be followed, this must be done in a way that does not interfere with efficient, reliable and cost-effective operation of LBNE.
- The File Catalog provides functionality that is very useful in a number of architectures of data management systems, which forms the basis for the corresponding requirement.
- Having the ability to always read raw data no matter when and how it was produced (including the format) is essential to meet requirements of Section 3.5.3 (data preservation).

3.13 Accepted Risks

- Having to implement a centralized system such as File Catalog leads to creation of a potential single point of failure. This will need to be countered with rigorous scalability analysis and implementation and adequate operational support.
- It is likely that at least part of the additional raw data replicas will be stored outside of the United States.
- Errors in decisions related to data retention will potentially result in necessity to reproduce large datasets.

3.14 Deviations

- Due to a hardware malfunction, a software bug or other such factors, a portion of raw data may fail basic validation. A special decision may be made to retire such data on a short time scale in order to conserve storage space.
- Once a particular body of data has been completely migrated to a new format, and properly validated, a special decision may be made to lift the "backward read capability" requirement for the software accessing the data, since the old format is effectively no longer required and the new software has full access to data anyhow.

4 Databases

4.1 Description

Databases are a crucial domain in LBNE, providing the foundation for handling experiment conditions and calibration data, metadata, file catalogs and a variety of information systems. A large part (but not necessarily all) information held in the experiment databases is typically derived from the experiment data itself (see Section 3 (Data)) and thus shares some of the same requirement concepts. However, due to specifics of data handling in this domain, databases are considered distinct enough that their requirements are stated separately.

4.2 Definitions

- **Database (DB):** A database is a collection of data, organized in a way that supports processes and activities requiring that data.
- **RDBMS:** Relational Database Management System.
- **Database server:** A database server is a computer program that provides database services to other computer programs or computers, as defined by the clientserver model. Most of the time the server runs on dedicated hardware, in which case the term "server" equally applies to that hardware as well.
- **Slave database:** A database in which information is synchronized to a master database, which is considered the authoritative source of data.
- **Master database:** A database which collects and stores information from various sources and serves as the source of data in the master/slave model. A replica residing in a slave may in turn become a master for a different system.
- **Database owner:** A collaborator or group of collaborators responsible for configuring, populating and maintaining a database and making it available to the Collaboration.
- **Collaboration database system:** The ensemble of all databases and their servers on which the collaboration relies.
- **Application-specific databases:** a database which is not required to be distributed because of its tight coupling to a specific application. Such databases are exempt from master/slave distribution mechanism as determined by consensus of the S&C Organization and its TAC.
- **noSQL:** A collective term referring to a large and diverse group of new database technologies which are built on principles different from RDBMS and may offer advantages in certain application areas

4.3 Issues

- In most physics experiments, RDBMS remains the prevailing technology in use, oftentimes as a legacy platform.
- There is no single RDBMS standard, there are multiple solutions and variation of the technologies as well as licensing options attached to it, from free to commercial.
- Commercial database solutions have the advantage of proven performance and well understood scalability, but also are often associated with considerable licensing fees, expensive support and implications for interface design.
- Data integrity requires having one definitive DB vs a few receiving same data simultaneously.
- Replication is one of core features necessary in many usage scenarios.
- DB backup/restore procedures are needed to protect crucial data.
- Performance and scalability problems are commonplace in large scale projects.
- The noSQL technology has taken root in industry, and should be evaluated for use in LBNE.

- Technology choices will have to be made sooner or later.
- There are multiple ways to access the data residing in the databases, and it needs to be done in a way optimal for each use case.

4.4 Requirements

4.4.1 Master Database

- Once a database is established for a specific type of data, there must be exactly one master DB for it, across the Collaboration.
- If technical implementation of the master database involves multiple servers sharing or copying data among themselves (such as in cases driven by High Availability considerations) such system must still preserve referential integrity of the data and have an interface presented to the slave DB servers (and other clients) consistent with that of a single master
- All master databases must be available for replication by any appropriately equipped participating computing facility, as long as it does not have undue performance and/or availability impact on the particular master DB.
- Design and implementation for new master databases shall be discussed within the S&C Organization and may be subject to recommendations by the S&C TAC if needed. Discussions shall take into consideration expertise of and ongoing support by the database owners, performance, scalability and long-term sustainability of the database technology as well as integration with collaboration-wide information systems.

4.4.2 Slave Databases

- A slave database must be replicated from a master and must not be otherwise modified.

4.4.3 Backup and restoration

- Any master database must be equipped with a backup mechanism or be otherwise reproducible in the event of a technical failure, hardware upgrade and other types data loss, on a time scale which does not significantly impact its ongoing use by the Collaboration.

4.4.4 Performance and scalability

- Specific metrics defining necessary availability, latency and scalability parameters for each database, such that their use is efficient and does not impede productivity of the Collaboration, will be established by the S&C Organization and its TAC.
- S&C Organization will be proactive in monitoring the performance of the Collaboration database systems, identifying bottlenecks and scalability limits, and planning data migration, equipment and platform upgrades and any other measures as is necessary to provide adequate level of database services to the collaboration.

4.4.5 Technology choices

- Database technologies available to the Collaboration will be considered on merits such as long term cost and performance characteristics, available expertise and support levels, available interfaces, scalability and other relevant factors.
- There will be no fixed requirement to either use or avoid newer technology such as noSQL vs RDBMS, and such determination will be made on case-by-case basis, taking into consideration existing and developing expertise and support by the database owners, cost estimates, performance, scalability and long-term sustainability of the database technology as well as integration with collaboration-wide information systems.
- Decisions pertinent to technology choices will be made based on discussions within the S&C Organization and if necessary, taking into account recommendations of TAC.

4.4.6 DB Access Methods

- The S&C Organization will make an concerted effort to insure that access to LBNE databases it not locked into a particular programming language or software package.
- Solutions coming from industry and open source community will be given preference over in-house development if other factors are equal.
- Where possible, direct DB access will be deprecated in favor of a Web service (or other comparable service) in order to take advantage of additional features such as security, data caching, more venues for scalability improvements, better resource utliziation via connection pooling and possibility to replace the database itself with a different engine, while keeping access methods intact.

4.5 Justifications

- Allowing replicas to receive updates from sources other than master DB will lead to loss of data integrity.
- It is essential to provide wide access to the data in the master DB, in order to leverage the data itself and the effort put into its maintenance.
- Optimal choices on specific master database location and technology can not be made without considering details of the application and the wider software and user base of the collaboration.
- Freedom of innovation and efficiency is gained by allowing flexibility in location and technology choices but this may conflict with collaboration wider goals of efficiency, security and consistency, and the right balance needs to be found.
- Master databases can benefit from geographical proximity to their source of data.
- Mature databases systems provide robust long-distance replication mechanisms which should not be re-invented by the collaboration.
- Opening a database connection to the world may lead to a sub-optimal security situation, and this can be mitigated by deploying a web service as an insulation layer.

4.6 Deviations

- Certain application-specific databases, such as a small database used in a small, limited use information system (e.g. a specific type of log files of no critical importance) will not be subject to replication requirements.
- Likewise, certain Web services require databases for their basic functionality, and such database contain information that is not specific or unique to LBNE data, and these cases will also be exempt from the replication requirement.
- In the interest of guaranteed availability and performance, replication requirements for online systems will be modified (e.g. ongoing automatic replication may be replaced with a managed copying of data, potentially in a different format, to systems that are exposed to users).

4.7 Use cases

- an application needs a straightforward way to read certain data in the conditions DB even when it is run off-line and does not have direct access to the Collaboration database systems. A read-only local replica of the data is created, which is immutable and will not be used to feed other systems
- since the performance of the master DB is critically important for it to be able to receive and process data at rates commensurate with incoming data, certain types of applications will be restricted to replicas of the data located in regional computing centers (as opposed to the master DB itself)
- after accumulating a large volume of data, an instance of a relational database system started having performance problems due to unforeseen scalability issue. Based on decisions by S&C, a noSQL service is established to serve monitoring data to Web applications and thus reduce the load on the principal database

5 Software

5.1 Description

The requirements in this section cover a variety of software used by the collaboration for producing its scientific results. Emphasis is made on architectural compliance and on managing the development cycle, validation in particular.

Requirements in this section cover the software required for collaborative work. Not included are engineering nor administrative software.

5.2 Definitions

- **LBNE Software:** The subset of simulation, reconstruction, analysis and production software which is written for the LBNE collaboration.
- **External Software:** Software on which the LBNE software relies. Also referred to as "externals" or "third party software".
- **Supported Platform:** A computing platform (OS and CPU architecture) which the S&C Infrastructure Group agrees to support.

- **User Friendly:** The ability to make use of something with only a reasonably minimal amount of effort by the intended qualified user.
- **Revision control:** (also known as version control and source control) is the management of changes to documents, computer programs, web sites, and other collections of information.

5.3 Issues

- Emphasis on well defined and optimal architectural patterns.
- Ensuring architectural compliance is difficult due to wide scope and complexity of LBNE software.
- Revision control.
- CASE and debugging tools.
- Release methodology.
- Collaborator and public access to repositories and commit privileges.
- LBNE software installation procedures.
- Defining supported platforms.
- Laptop or workstation applications vs. production (use case important here).
- Adoption of external tools, review process.
- Collaborator produced software, review process.
- Software documentation.
- Proprietary software acceptance.
- Distribution of software.

5.4 Software Architecture and Architectural Compliance Requirements

- All components of production software will be created according to appropriately chosen architectural patterns, aimed at maximizing maintainability and adaptability.
- Basic Architectural Compliance will be ensured in the processing of code review and subsequent testing (see Section 5.8), taking into account implementation of components, interfaces and behaviors.

5.5 External Software Requirements

- **Provenance:** the S&C Organization shall assure that all third party software available in source form can be compiled from pristine upstream releases and any applied patches are taken from trusted sources.
- **Source code archival:** a copy of the source code used to build external software for official releases will be archived in order to allow for reproducing results at a later time
- **reliance on third party binaries:** reliance on external software which is available only in binary form will be kept to an absolute minimum and considered on a case-by-case basis.

5.6 Requirements for Supported Platforms and Versions

- The S&C Organization shall develop a policy for supported platforms, based on the needs of the Collaboration. This will include the mechanisms by which platforms are selected for support, retired from the supported status and what levels of support and methods for its delivery will be provided
- the S&C Organization determines the supported versions of software, taking into account the computing needs of the Collaboration

5.7 Documentation Requirements

- All important aspects of the LBNE software design, development, commissioning and operation will be documented using a documentation system selected by S&C Organization (see Section 10.3).

5.8 General Code Management Requirements

- The S&C Organization shall develop and document a set of coding standards to be followed by all experiment software.
- All production software and any software of similar importance as determined by the S&C Organization and the TAC will be subject to code review.
- Code reviews will aim to satisfy the requirements of Architectural Compliance and coding standards.
- Code review feedback and results will be documented according to general documentation requirements, in order to facilitate access to this information and its usefulness for making corrections.
- All LBNE software source code shall be maintained in a version control system.
- Backup and restore procedures must be established in counter possible loss of source code due to technical failures.
- The S&C Organization will establish which debugging, CASE and other automated tools will be acceptable for use in various sectors of software development in LBNE.
- Any code used directly in LBNE production software or representing a dependency for it for must be preserved for the lifetime of the Collaboration, in order to guarantee its ability to reproduce all and any of its results at any given time.
- A review process will be implemented to determine which parts of LBNE software are no longer required and/or used, and in cases like these decisions will be made whether and how to decommission and remove such software in a way that is not disruptive to LBNE.
- A release mechanism will be put in place, including defined release periods for LBNE software and externals.
- A policy on long term software maintenance and knowledge management will be put in place in order to address eventual rotation of personnel and departure of software developers from the Collaboration.

- Software build and installation from source: an installation mechanism based primarily on the source code will be provided for the LBNE software and externals which can be run on compatible platforms, in an automated fashion by collaborators with a limited but reasonable amount of expertise.
- There will be an established and documented process which developers of new LBNE software packages should follow in order to integrate the package into the greater code base.
- Provision of software binaries: the S&C Organization must provide packages in a form ready to use (e.g. executable binaries, libraries, properly placed Python modules or scripts) and a mechanism to install them in a user friendly manner.

5.9 Validation, Unit Testing and Continuous Integration

5.9.1 Definitions

- **validation:** in this context, we define it as software quality control
- **unit test:** a method by which individual units of source code, sets of one or more computer program modules together with associated control data, usage procedures, and operating procedures are tested to determine if they are fit for use
- **continuous integration** frequent updates of the code in the common repository, combined with running unit tests in an on-going manner
- **test harness:** a system supporting the creation and execution of unit tests

5.9.2 General Validation Requirements

- Testing procedures will be designed in a way which allows them to be run in fully automatic as well as more interactive, user-controlled mode.
- All of the software used in LBNE must be valid on all supported platforms.
- For a given unit version or release, or an external package, to be declared valid, all of its components must compile in the LBNE environment.
- For an application or program to be declared valid after build procedures are complete, there must be no errors introduced in the process.
- Methods must be provided to reproduce results from prior runs of the same program.
- Regression test failures must be thoroughly investigated and result in one of the following:
 - Identification of bugs introduced since the last successful regression test.
 - Expected new behavior due to code changes, which must result in updating the test suit.
 - New features of the environment, which need to be accounted for, documented and taken into account while updating the test suite.
- Software will be tested (where necessary) for compliance and compatibility with evolution of data schemas and formats.

5.9.3 Unit Testing and Continuous Integration

- Every LBNE software package must provide unit testing capability.
- There will be a Continuous Integration System (CIS), which will cover all supported platforms and include necessary test harness capability.
- The CIS will rely on the unit tests coming with the packages it covers.
- Resolution of issues (e.g. bugs) must be confirmed through unit testing.
- The CIS and the unit test upon which it relies will include checks for both basic validity of the code (e.g. it doesn't crash) and for the physics content of the code (e.g. a particular observable being calculated matches the expected value according to set criteria).
- There will be a link between the CIS and if necessary, individual unit tests, and the Issue Tracking System deployed by the Collaboration (see Section 10.3).
- Full diagnostic information generated by CIS will be stored in a way that is accessible to the members of the Collaboration, the full retention policy will be determined by the S&C Organization at a later time.

5.10 Software Distribution and Configuration

5.10.1 Description

The software distribution is the mechanism by which the software intended to run on a particular host or facility is being made available for this purpose. This is a non-exclusive list of methods that can be employed for this:

- Local installation from source code, as described above, resulting in placement of binaries and other component in a local file system accessible to the Worker Nodes, desktops and other computing devices and services at a particular location.
- Building the binaries from source at a central facility, for all supported platforms, and subsequent managed "push" of the software pre-built in this manner to remote sites.
- Using network file systems such as CVMFS and others, to make the software accessible to the target computer without having completed local deployment

5.10.2 Requirements

- LBNE S&C will implement a system where the software is built in a centralized manner and then distributed to remote computing centers (this does not imply that it's the only available or endorsed method).
- There will be a common configuration mechanism to ensure proper functioning of the software on each site.
- By utilizing networks to distribute data on demand, LBNE S&C will provide a mechanism for Worker Nodes to gain access to software binaries even in absence of local installation at a particular site.

5.11 Use Cases

- During a code review, the S&C reviewers identified hardcoded values for certain geometrical parameters. Findings are documented in later used to rectify the problem by creating an appropriate interface to the geometry database.
- In the process of a nightly build, a recently added module failed to compile. According to the S&C policy, the entire release is not considered valid and cannot be used in production.
- A request is received from one of the users to provide help in installing LBNE software on a new version of Linux they installed on a desktop system. Since this platform is not officially supported, action is deferred and assistance provided as situation permits on best effort basis.

6 Distributed Computing

6.1 Definitions

- **Grid computing:** technology that allows a collection of computer resources from multiple locations to reach a common goal.
- **Cloud computing:** distributed computing technology with emphasis on network connectivity, virtualization and elasticity of resource allocation.
- **Workload Management System (WMS):** a system that enables automated placement of computational payload jobs submitted by its users on distributed resources, using the underlying Grid layer, and makes subsequent record keeping, accounting, elements of data management and general monitoring available to the user.
- **Monitoring:** in this context, it represents a system which stores the state of individual jobs, groups of jobs, data transmission and other objects vital for Workload Management, while providing appropriate interfaces to end-users and operators of a WMS.

6.2 Grid and Cloud Computing Capabilities

6.2.1 Description

This section describes the basic design principles and approaches in the area of organizing the distributed computing resources being utilized by the Collaboration.

Grid and Distributed Data technology has been embraced by virtually every major scientific project in recent memory, for reasons of scalability and efficient access to resources. LBNE needs to closely examine its approach to this domain and create a strategy for integration of Grid and Cloud Computing capabilities into its infrastructure base.

6.2.2 Issues

- Most efficient utilization of all resources available to the Collaboration, which includes both efficient use of all available hardware, and the ability of LBNE to access the resource in a transparent, agile and efficient manner.

- Hedging against significant uncertainties inherent in estimating and planning resource allocation over the next 10 years, against the backdrop of the quickly evolving technology landscape.
- Grid Sites may have a wide range of capabilities, interfaces and other configuration parameters that need to be accounted for in order for these Sites to be properly interfaced and integrated.

6.2.3 Requirements

- LBNE S&C Organization shall implement a widely distributed computing infrastructure, featuring a network of federated resources (including Grid- and Cloud-based resources).
- The distributed resources can include processing facilities, storage, network hubs and their combinations (e.g. Grid sites with large available storage capacity).
- Necessary tools and procedures will be provided, for streamlined incorporation of new facilities as they become available, efficient use of opportunistic resources and leveraging the effort and expertise across all of the Collaboration in order to provide adequate operational support with a minimum of manpower.
- Details of site capabilities, interfaces, configuration and other crucial information will be gathered and preserved as documents and database records, and made available through appropriate information systems (often termed Grid Information System).

6.3 Workload management

6.3.1 Description

Workload Management Systems serve as a force multiplier for the user of Grid technology. They insulate individual users from specific configuration details and certain failure modes of Grid sites, provide substantial automation in managing the user's computational payload on the Grid. A good WMS also must provide adequate monitoring capabilities (down to the job level) which serve as a valuable debugging tool, and represent an essential toolkit for the operational support teams.

A WMS is also a valuable tool for managing software deployment, keep proper accounting information regarding releases and document configuration (e.g. release) of the software used for a specific production run, thus satisfying one of core requirements related to processed data in LBNE.

6.3.2 Issues

- Due to a very large number of individual computational payload jobs, wide variations of their type, requirements, environment and other vital characteristics, having access to Grid facilities in itself does not guarantee acceptable degree of efficient utilization and availability of resources, and/or functioning Operational Support.
- Agility in resource allocation and sufficient monitoring capability is not included in the typical basic Grid middleware stack, while being essential for efficient operation of the experiment.
- A robust and user-friendly web information system, allowing individual researchers and working groups to track the execution of, and identifying potential issues related to their payload submitted to the Grid, is a significant force multiplier, leading to increase in productivity and validity of results obtained.

6.3.3 Requirements

- LBNE will implement a Workload Management System (WMS) for resource management and brokerage procedures, which will govern distribution of most types of computational workload in LBNE (e.g. production jobs, group analysis etc) across resources available to the Collaboration.
- The LBNE WMS will be capable of keeping precise record of the software configuration used for each and every job deployed on the Grid, including, among other things, the LBNE Software Release information.
- The LBNE WMS will be capable of quickly suspending participating sites due to outages, network congestion or potential security issues.
- The WMS will be augmented with a Workflow Management layer, which will help create and manage large groups of Grid tasks supporting the scientific workflows.
- LBNE will also implement a monitoring system to allow efficient operation of its WMS, by helping ascertain correct execution of Grid jobs, accounting of resource utilization, identification and debugging of failure modes etc.
- The WMS Monitoring System will keep status records for individual jobs and their groups, information related to data I/O and transmission, and include the "big picture" performance monitoring data for entire Grid sites utilized by LBNE.
- The WMS Monitoring System will have interfaces conducive to integration with both Web UI for users and operators, and with automated systems needing the WMS data.

6.4 Technology trends in Distributed Computing

6.4.1 Description

Technology is progressing at a brisk pace and it's virtually impossible to cover its various implications for LBNE Distributed Computing within the scope of this document. We shall itemize, however, a few directions that we'll need to explore in the next few years, such as HPC, multi-core and multi-threaded processing, GPUs etc.

6.4.2 Issues

- High-Performance Computing (HPC) platforms have very different characteristics from the High-Throughput (HT) systems commonly used in HEP.
- Multi-core has become a standard feature of the hardware utilized at data centers everywhere, and its efficient utilization remains a challenge
- GPUs continue to offer exceptional performance in a few application domains, while requiring specialized software development approaches, tools and architectures.

6.4.3 Requirements

- LBNE shall be proactive in its assessment of multi-core and other multiprocessor technology in order to assure efficient resource utilization and leveraging available technologies for better performance of its physics tools software.

- S&C Organization will work in cooperation with the LBNE Physics Tools Group to explore potential use of GPUs and HPC facilities for LBNE needs.
- S&C TAC will be tasked with making recommendations for multi-core implementations of LBNE software and specific approaches and standards that must be followed in this area.

6.5 Use Cases

- A LBNE member institution prefers to be able to host and process significant parts of the data, and to run some new type of analyses locally for expediency and efficiency, while still retaining the possibility of using larger opportunistic resources available to LBNE as a whole, if necessary. This is made possible by using the LBNE Workload Management System and its distributed data handling mechanism.
- While running managed production at FNAL, it is discovered that the deadline for delivering results to an important conference cannot be met assuming only local resources are available. At the same time, certain member institutions possess extra computing power that can be used to mitigate the shortage. Additional Grid jobs are dispatched, thus filling the need for more resources.
- A small team of researchers are running an important MC simulation on their very limited local resources, and need to accumulate a lot more statistics than can be reasonably expected with what is available to them. They take advantage of facilities at FNAL, which is made possible by transparency of access and data handling, and user-friendly Grid interfaces.
- A user observes that a large portion of her jobs started failing according to the monitoring data coming from the WMS. Using the job-level monitoring facilities provided by LBNE, she's able to pinpoint and fix the problem.
- There is a new and intermittent problem with network connectivity to one of LBNE sites, resulting in excessive data transmission times and inefficient resource utilization. This condition is recognized by LBNE Operations Team and the site is suspended for WMS purposes while the problem is worked on.

7 Geometry

7.1 Description

Managing geometry description and model presents a challenge to most experiments, and LBNE is no exception. Ideally, there is "one source" geometry description which feeds both Monte Carlo simulations and reconstruction of data, and potentially takes advantage of interfaces to CAD systems etc. The obvious benefits include assurance of the integrity of the geometry data, and leveraging visualization systems across various sectors of the experiment. In practice, such complete integration is hard to achieve. Nevertheless, we identify a few requirements in this area that should be met.

7.2 Definitions

- geometry description :: a collection of information sufficient for creation of the geometry model of the detector, for the purposes of either simulation, reconstruction or both

- geometry model :: a collection of data structures plus the code to manipulate these, which together provide the functionality required by the application (simulation etc, cf. the geometry model in GEANT4)

7.3 Issues

- a variety of domains utilizing geometry data, such as simulation, reconstruction, packages specific to the Far and Near Detectors, beamline description, overburden maps, interaction packages
- legacy code and ad hoc solutions already utilized in these sectors
- compelling advantages of a single, unified system, which is however difficult to achieve
- necessity to implement replication and symmetry exploitation in the geometry description and model, and a well structured, single source geometry code/description
- validation mechanisms: visual, overlaps detection, tracking tests
- interfaces that need to be built to existing and future tracking and visualization packages

7.4 Requirements

- LBNE shall develop a central geometry management system from which specific applications may derive their geometry information
- such system must provide geometry information for near, far and prototype detectors, for beamlines, for overburden and other geometry information used for collaboration results
- the system must support reliable storage and revision control of any set of geometry information

8 Visualization

8.1 Description

This section describes requirements related to visualization tools used to produce graphics as required by LBNE. CAD and other engineering systems are not included here. Visualization in this context is the process of displaying both the structure of various components of the LBNE apparatus, and optionally also the event data. A special case of combined visual representation of both is often referred to as "event display".

8.2 Issues

- importance of visualization tools for many types of research, development and data analysis work in LBNE
- variety of legacy geometry descriptions, split across the detector subsystems, and variety of models that need to be interfaced
- the liquid Argon TPC is a precision tracking apparatus and by its very nature a "visual" detector, thus having robust and versatile ways to display data coming the TPC will leverage its capabilities
- the role of visualization tools in the development of tracking algorithms of LBNE is especially important

8.3 Requirements

- the S&C Organization shall provide applications and supporting tools to visualize the LBNE geometry and event data
- visualization toolkit available to LBNE will cover a range of functionalities, including interactive event displays, and capability to produce publication-grade graphics
- provisions will be made to accommodate specific needs of simulation work, such as displaying "Monte Carlo truth" information
- the suite of visualization tools will cover the needs of expert users (such as during event scanning studies) as well as general needs of the collaborators who do not necessarily possess significant programming expertise

9 Networks

9.1 Description

The network infrastructure of LBNE is important for two main reasons. One, traditional reason is that LBNE will depend on the networks for distribution of its data to storage facilities and computing centers located at its member institutions, as well as distribution and management of its computational workload on the Grid. This will require a focused effort on design, deployment and continuous monitoring and maintenance of the LBNE networking infrastructure. The other and somewhat more unusual reason that networks must be a focus area in LBNE is due to the fact that the Far Detector is located indeed very far from the beam, the target and tape storage at FNAL.

9.2 Far Site Connectivity

9.2.1 Issues

- Beam on target signal containing timing information about the beam, to be used by the online systems
- Transmission of raw data out of the detector and to the storage facility, tentatively located at FNAL, which also ties into the topic of buffering the data to compensate for occasional problems with the network connectivity and performance

9.2.2 Requirements

- a R&D effort will be undertaken to create the necessary infrastructure, including software, for the "Beam on Target" signal
- the S&C Organization will conduct research and establish performance parameters for the network connection between the "Far Site" detector and principal LBNE data storage centers (e.g. FNAL)

9.3 LBNE Data Network

9.3.1 Issues

- LBNE will utilize a global and complex network to distribute, collect and otherwise manage its data across multiple data centers around the globe.
- Individual elements of the networks utilized by LBNE will have vastly varying performance and reliability characteristics, subject to change in time.
- In order to utilize resources in an efficient manner, LBNE will need to be agile in its utilization of networks.

9.3.2 Requirements

- A monitoring system will be put in place to continuously gauge the bandwidth, availability and other vital parameters of network connection established among participating LBNE sites.
- Network metrics shall be collected and stored for analysis, and used for optimizing data distribution algorithms and policies.

9.3.3 Use Cases

- A data processing campaign is under way, and input data needs to be shipped to a few sites. The monitoring system deployed by LBNE collects network performance metrics that indicate that it won't be optimal to transmit data to a few sites due to insufficient bandwidth, and these sites are subsequently used mainly for Monte Carlo studies. At a later point, when the network problems are resolved, primary data transmission is resumed. The net result is better efficiency of the overall system.
- A sudden connectivity glitch between two remote sites is detected by an automated system, and a notification is sent to LBNE Operations Support.

9.3.4 Alternatives

In case a centralized network monitoring system is not created, monitoring the network functionality and performance can be delegated to personnel located at individual sites.

10 Collaborative Tools

10.0.5 Definitions

- **Content Management System (CMS):** a computer program that allows publishing, editing and modifying content as well as maintenance from a central interface.

10.1 Description

Collaborative Tools for LBNE will include the following:

- Issue Tracking Systems - these can potentially cover a very wide range of issues pertaining to software, from documenting progress on detection and elimination of bugs to issues of physics performance and managing scientific workflows.
- Document Management Systems.
- Content Management Systems, purpose-built information systems etc.
- Teleconferencing tools.
- Meeting and conference management tools.

10.2 Issues

- An efficient, user-friendly system for issue tracking is absolutely crucial for LBNE to maintain a functioning Software and Computing infrastructure.
- Keeping up-to-date, easily accessible documentation is absolutely crucial for adequate rate of progress of LBNE Software and Computing.
- Phone and video conferencing tools are powerful project management tools.
- LBNE is a large international organization and picking the tools most widely utilized in the community will lead to shorter learning curves and added efficiencies.
- In managing various kinds of data and metadata in LBNE, it may become necessary to develop and deploy information systems accessing databases and other data sources, which go beyond single application domain and can be used by an entire Collaboration.

10.3 Requirements

- A single issue tracking system will be put in place to cover all aspects of work in the S&C Organization, the specific type of the system to be determined by S&C and its TAC on a time scale of a few months.
- The system used for issue tracking must have a proper network interface such that it can be used in non-interactive and automated applications and integration projects (e.g. RESTful interface).
- S&C Organization will evaluate various CMS available for its use, collect feedback from its members and make a determination of what system shall be used; in the interim, legacy systems will continue to be used for immediate needs of the Collaboration.
- Chosen CMS will have to include auth/auth and granular access control functionality, revision control and formatting tools.
- No new CMS will be developed in-house.
- Teleconferencing tools will be chosen to provide reliable connectivity for at least 100 concurrent callers, and a functional set of common conference management tools (e.g. ability to mute users, manage the duration of the conference etc).

- A document management system will be used, which will include authorship management, revision controls and review/approval functionality.
- A meeting/conference scheduling and organization tool will be identified based on preferences of the Collaboration, functionality and ease of use.
- Information systems and their Web interfaces will be based on existing and proven technologies e.g. web frameworks, as opposed to creating in-house solutions for DB access and Web UI.

10.4 Justifications

- Based on experience, rapid one-off development of an ad hoc information system using custom code, or in-house design of a content-management system are tempting, but it's not sustainable in the long term and does not allow for re-using industry standard and/or open source packages and solutions.
- Unreliable phone conferencing system invariably results in lost productivity.

10.5 Use cases

- A user notices malfunction in one of the jobs, via the available monitoring tools, and creates an entry in the issue tracking system, while attaching a link to the item information to the job object in the monitoring system. The entry in the issue tracking system proves useful for other users in a similar situation.
- There is a need to make a certain kind of calibration information easily accessible to the users. While designing a system to satisfy this need, a decision is made to use an industry-standard, proven solution (web framework) which enjoys a lot of community support.
- A Collaboration note on a software topic is being published after a review which is conducted and managed using the LBNE document management system.
- A group of users is using the LBNE issue tracking system to manage a particular complex computational workflow they need to complete for a deliverable.

11 Cybersecurity

11.1 Description

Attacks on scientific computing facilities by malevolent agents are a reality that needs to be countered. Due to users being lax in fully implementing security protocols, improper site configuration and other factors, these attacks are successful at times.

11.2 Issues

- It is impossible to prevent malevolent entities from attacking the LBNE computing facilities, such attacks are commonplace.
- There are proven methodologies, policies and approaches as well as rules put in place by the agencies, that minimize the success rate of such attacks, and these policies and rules must be fully followed.

11.3 Requirements

- LBNE S&C will establish close cooperation with Cybersecurity personnel at participating sites, and in particular, these sites will be required to report all security related incident to S&C on a monthly basis or more frequently in a crisis situation.
- All rules and policies set forth by DOE and other authoritative sources must be implemented.
- Protocols will be established to stop and mitigate attacks.
- LBNE S&C Organization will periodically review the configuration and security features of its member sites.
- Creation and use of "production accounts", i.e. user identities associated with automated services, agents and other such type of computer processes, will be kept to absolute minimum, reviewed on case-by-case basis and a subject to strict information protection.
- There will be a blacklisting capability in LBNE, at the very least covering its Workload Management System but also other systems as needed, which will make it possible to block compromised or abused accounts on a short notice.

11.4 Use Cases

- Unusual activity is detected on one of LBNE distributed sites. Local personnel determines that this is a case of a hijacked account used to mount a DOS attack. Corresponding processes are located and eliminated, while the account in question is blacklisted until further investigation and corrective measures.
- LBNE interacts with organizations like the Open Science Grid in order to conduct an internal review of its Cybersecurity protocols and implement best practices going forward.

References

- [1] LBNE Conceptual Design Report
- [2] LBNE Software and Computing and the Physics Tools Plan