

# The Fabric for Frontier Experiments Project at Fermilab

Michael Kirby<sup>1</sup>

<sup>1</sup>Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL, USA 60510

E-mail: [kirby@fnal.gov](mailto:kirby@fnal.gov)

## Abstract.

The FabrIc for Frontier Experiments (FIFE) project is a new far-reaching, major-impact initiative within the Fermilab Scientific Computing Division to drive the future of computing services for experiments at FNAL and elsewhere. It is a collaborative effort between computing professionals and experiment scientists to produce an end-to-end, fully integrated set of services for computing on the grid and clouds, managing data, accessing databases, and collaborating within experiments. FIFE includes 1) easy to use job submission services for processing physics tasks on the Open Science Grid and elsewhere; 2) an extensive data management system for managing local and remote caches, cataloging, querying, moving, and tracking the use of data; 3) custom and generic database applications for calibrations, beam information, and other purposes; 4) collaboration tools including an electronic log book, speakers bureau database, and experiment membership database. All of these aspects will be discussed in detail. FIFE sets the direction of computing at Fermilab experiments now and in the future, and therefore is a major driver in the design of computing services world wide.

## 1. Introduction

As Fermilab becomes the world leader in Intensity Frontier (IF) particle physics research, the computing requirements of IF experiments will increase by an order of magnitude compared to current on-site resources. To enable the IF experiments to explore new physics utilizing intense beams of neutrinos, muons, kaons, and nuclei, IF experiments with the help of Scientific Computing Division (SCD) will need to design larger and more complicated computing models than previous IF experiments. Historically, IF experiment computing models at Fermilab have been limited in the ability to utilize off-site computing. But it is apparent that the ability to rapidly and efficiently utilize distributed computing resources will play a critical role in analyzing multi-petabyte size datasets from next generation IF experiments. The transition to distributed computing has been at times difficult for experiments and not without significant challenges. At the same time, Fermilab SCD and the Open Science Grid have created world-leading computer facilities and tools for distributed computing. The FabrIc for Frontier Experiments (FIFE) project aims to incorporate existing tools into an integrated framework to improve the design and implementation of these resources for IF experiments. By integrating these facilities and tools into the computing model of Intensity Frontier experiments, FIFE will enable them to achieve their physics goals and perhaps beyond.

## 2. Computing Needs at Fermilab

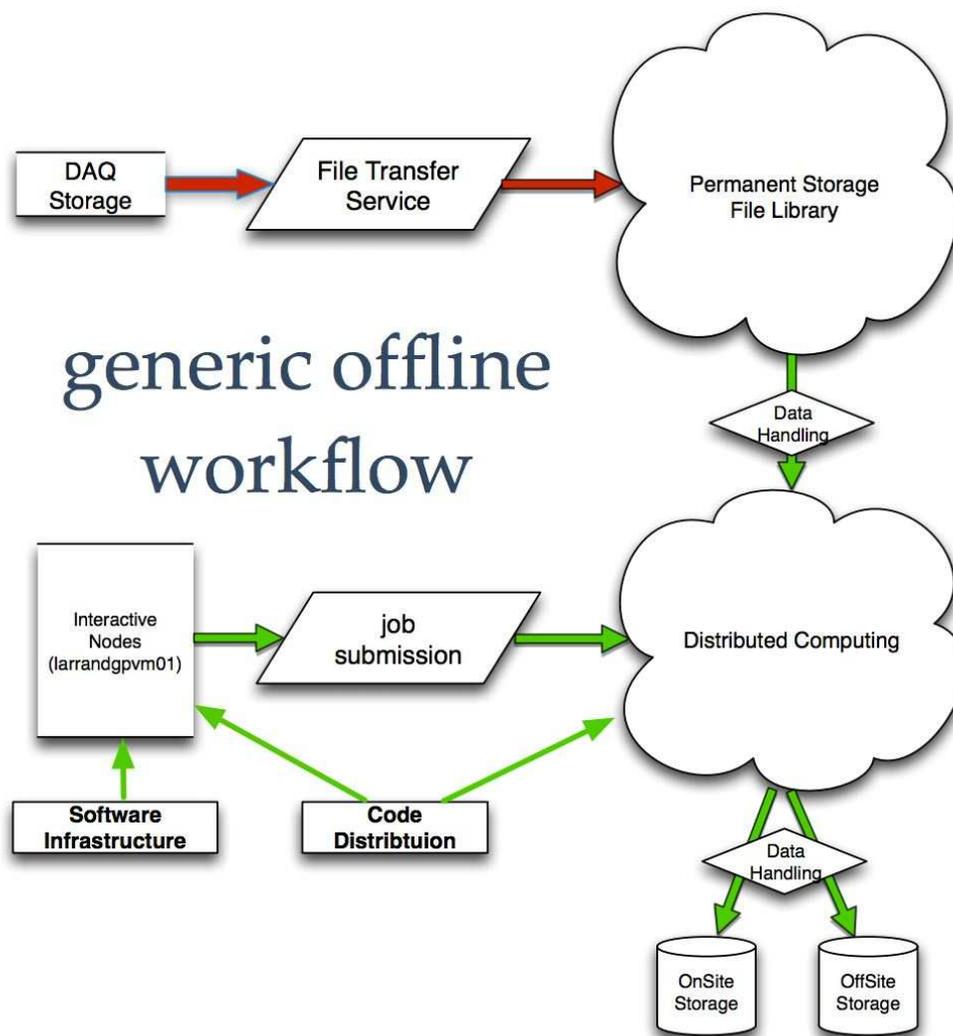
In the next decade, the extensive Fermilab scientific program exploring questions at the Energy, Intensity, and Cosmic frontiers will pose significant requirements for computing resources. The Intensity Frontier program at Fermilab includes neutrino experiments (e.g. MINOS+, MINER $\nu$ A, NO $\nu$ A,  $\mu$ BooNE, and LBNE), muon experiments (Muon  $g - 2$  and  $\mu 2e$ ), along with proposals for experiments studying rare decays. Cosmic Frontier experiments are being explored with the Dark Energy Survey and DarkSide-50 and proposals for second generation dark matter and dark energy experiments. The expectation is for the active operation of 8 experiments in FY16 and to have deployed all of the computing resources need to perform numerous tasks for those experiments. They will require resources for beam simulations, detector design studies, production event reconstruction, physics event generation, and detector response simulation. With the advent of larger datasets, the importance of effective offline computing design has come to the forefront for all aspects of particle physics. While each new experiment brings new challenges in computing, with some analyses that are I/O limited, some CPU limited, it is apparent that distributed computing is a necessity for IF experiments to be successful.

Unfortunately for many of the new experiments at Fermilab, the limited number of collaborators on the experiment translates to the inability to develop computing infrastructure within the experiment. For both the ATLAS and CMS experiments, the large number of collaborators  $O(>3000)$  on these experiments allows for a significant amount of person power to be dedicated to the development of computing infrastructure to be developed within the experiment. For smaller experiments such as  $\mu$ BooNE and Dark Energy Survey having only 100 collaborators means that dedicating resources to architecting and implemented computing infrastructure for distributed computing is difficult. But the failure to couple world-class computing resources to this new experimental data would significantly limit the success of these experiments and hinder the exploration of ever expanding frontiers.

In order to enable distributed computing for IF experiments, the Scientific Computing Division has started the FabrIc for Frontier Experiments (FIFE) project with the aim of providing a consistent OSG focused framework for experiments to utilize. Where appropriate, FIFE will reuse current solutions from other OSG communities. FIFE will consult with experiments on the applicability of supported services, aid in implementation, and develop new solutions for the unique needs of experiments. The availability of a well designed, efficient, and powerful computing model is critical to success of the new experiments and FIFE will be at the forefront of this process for Fermilab. The vision of location-agnostic computing models will allow Fermilab and it's collaborating institutions to maintain its position as a world leader at all three Frontiers and fully utilize all available resources.

The overarching goals of the FIFE project are:

- enable world class science through high performance computing
- provide an integrated framework that allows researchers to focus on analysis algorithms with little focus on computing infrastructure
- consult with experiments to improve offline computing models where appropriate based on the FIFE framework and tools
- enable optimal and convenient use of computing resources at Fermilab, outside research institutions, other national laboratories, and commercial computing providers
- provide solutions that are modular in design so that experiments may utilize only solutions they require
- provide consistent interfaces so that underlying solution changes do not modify user interaction
- become a world leader in the development, adoption, and deployment of distributed computing tools



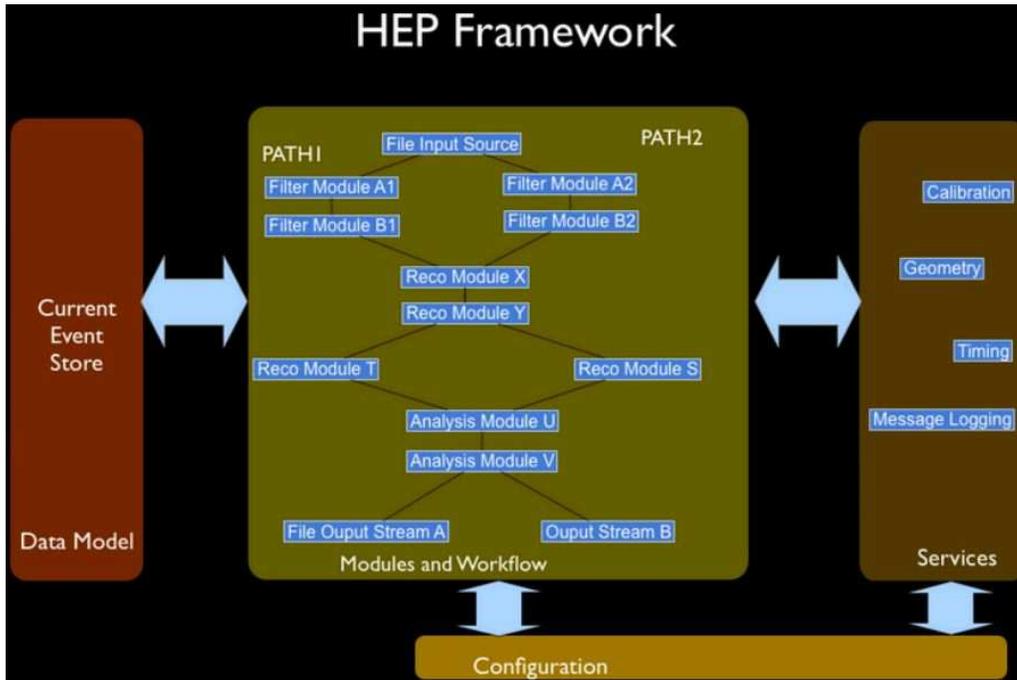
**Figure 1.** Simplified model of the Fermilab experiment offline computing process.

### 3. Generalized Solution for Data Processing

While not universal in addressing all computing models, the FIFE architecture design has been based upon a simplified model for experimental computing requirements. The simplified model is built upon the following components: software infrastructure, code distribution, data handling, database infrastructure, job submission and monitoring. A diagram of the simplified model is shown in Figure 1. The details of the current solution for each component are described below.

#### 3.1. Software Infrastructure

In order to provide physicists with a software infrastructure for analysis and reconstruction development, the Scientific Computing Division has deployed the art[1] software framework. The purpose of the art framework is to provide a well defined environment designed to encourage best practices in algorithm design and coding implementation. The framework provides tools for multiple processing paths, provenance of algorithm configuration, and metadata production for the output of processing. The design of modules, workflows, and services were specifically



**Figure 2.** A diagram of the art framework showing multiple data processing paths and the interaction between modules and the event data record.

designed in order to assure that output are consistent and reproducible with minimal effort. art has also been integrated with ROOT, GEANT4[2], GENIE[3] and several other standard high-energy particles physics software packages. Figure 2 shows a diagram of the art framework where multiple paths contain several modules and the interactions with services and event data. The art framework is currently in use by the ArgoNeuT, Muon  $g - 2$ ,  $\mu 2e$ , LBNE,  $\mu$ BooNE, NO $\nu$ A, and DarkSide 50 experiments.

Along with providing a software framework, the FIFE project is working to provide resources to build and distribute experimental software. Typically, an experiment will build the shared object libraries of their software each night for distribution to analyzers. With hundreds of thousands of lines of code and several flavors of libraries to produce, the process of building an experiment’s software can become long enough to disrupt the distribution. In order to address this, the FIFE project is developing a service whereby dedicated hardware which has been optimized for quickly building an experiment’s software would be time shared between all FIFE experiments. When combined with the OSG OASIS[4] server utilizing the CERN Virtual Machine File System[5], this allows for minimal disruption in code distribution to the experiments.

### 3.2. Job Submission

One of the most important aspects of the experimental interface for accessing grid resources are the tools utilized to submit jobs. The FIFE project is working to modify the existing job submission tools currently in use by the Intensity Frontier experiments to be more focused on the OSG. This transition will most importantly change to a client-server model where lightweight client software that can be installed at several platforms will interface over a REST-ful protocol to the job submission server to schedule processing on Fermilab local and remote computing elements. The FIFE job submission tools rely on the GlideinWMS system on the back end to

present distributed computing resources as a single pool of worker nodes. This system provides a consistent running environment for jobs and allow for quick integration of new remote resources. As well, the job submission services handles the distribution and renewal of grid certificates to ensure that data transfers to and from job sections is never interrupted. The interface with data handling, code distribution, and workflows has also been incorporated within the job submissions tools. Most importantly, the FIFE architecture has been designed that this service will maintain a consistent interface while the underlying solution may be modified so that experimental scripts and knowledge will not be lost during such a transition.

### *3.3. Data Handling*

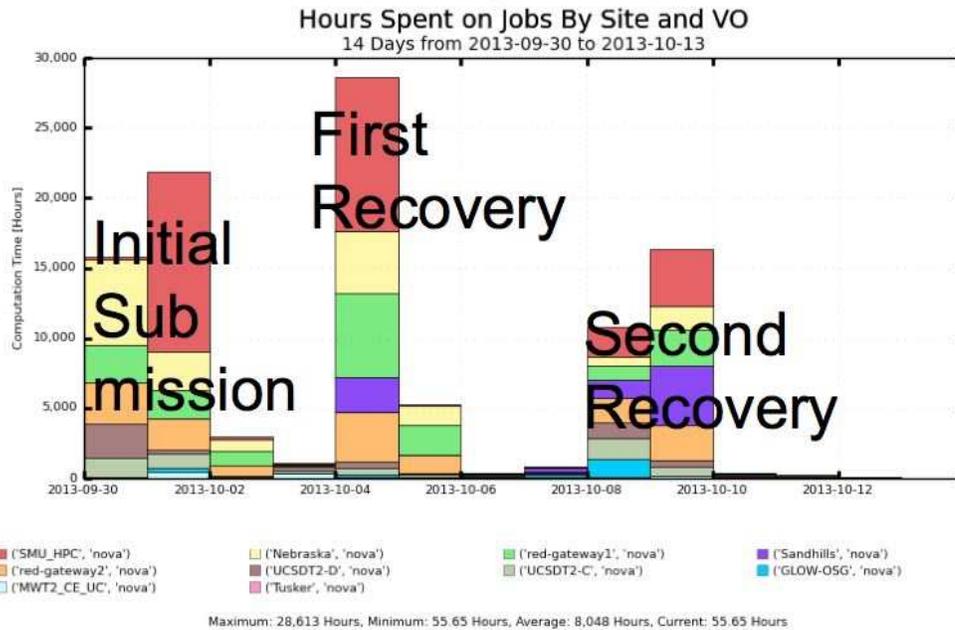
Within the offline processing model for FIFE, one of the largest challenges for analysis is the delivery of data to computing resources. There are three essential tasks in data delivery within the FIFE architecture: storage element infrastructure, file catalogs, and transfer services. FIFE is actively deploying large scale (many PetaByte) storage elements using two different protocols: BlueArc hardware via NFS and dCache. The Serial Access via Metadata (SAM) service is a robust and mature file catalog, delivery and tracking system developed at Fermilab based on file metadata while being storage element independent. The SAM service has recently made a transition to a web based interface (SAMWeb) and allows for greater integration into OSG operations. The Intensity Frontier Data Handling Client (IFDHC) service is designed to provide access to all Fermilab storage elements (BlueArc, dCache, and Enstore tape storage) through a single interface. The IFDHC is focused on making sure that resource brokering keeps any storage element from becoming overloaded while minimizing worker idle time waiting for resource tokens. While currently being explored by other projects, FIFE does not currently have immediate plans for delivering computing jobs to the location of stored data.

## **4. NO $\nu$ A Integration**

Recently the NO $\nu$ A experiment integrated the FIFE infrastructure into their offline computing model in order to increase the cosmic ray background simulation samples by an order of magnitude for a total of one million events. Utilizing all of the FIFE elements (code distribution, job submission, and data handling), NO $\nu$ A was able to generate samples at six remote sites: University of Nebraska, Southern Methodist University, University of Wisconsin, University of Chicago, and University of California - San Diego. A monitoring plot of the NO $\nu$ A CPU hours at the remote OSG sites is shown in Figure 3. While the focus of the integration effort was the utilization of opportunistic OSG CPU hours, an additional proof-of-principle submission was performed using virtual machines on FermiCloud that were launched on demand. With the successful integration of NO $\nu$ A into the FIFE architecture, the expectation is that the FIFE project will be able to quickly integrate additional experiments.

## **5. Conclusions**

With the advent of larger datasets, the importance of effective offline computing design has come to the forefront for all aspects of particle physics. While each new experiment brings new challenges in computing, with some analyses that are I/O limited, some CPU limited, it is apparent that distributed computing is a necessity for IF experiments to be successful. The coupling of world leading computing resources with experimental data allows for the exploration of ever expanding frontiers. The availability of a well designed, efficient, and powerful computing model is critical to success and FIFE will be at the forefront of this process for Fermilab. The vision of location-agnostic computing models will allow Fermilab and its collaborating institutions to maintain its position as a world leader at the Intensity Frontier and fully utilize all available resources. As well FIFE will define the scope, goals, and timescale for scientific computing in the coming decade. All the while, FIFE will enable the exploration of new frontiers



**Figure 3.** A plot of the  $\text{NO}\nu\text{A}$  CPU hours on the Open Science Grid during the last week of September 2013 and the first week of October 2013. The three peaks are from the initial submissions of jobs and then two secondary submissions to complete the cosmic simulation samples.

in physics that will make some of the most significant measurements and discoveries in particle physics.

### Acknowledgments

The author wishes to thank the  $\text{NO}\nu\text{A}$  experiment and all FIFE project members for extensive effort in initial integration of the  $\text{NO}\nu\text{A}$  offline computing. Fermilab is operated by the Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.

### References

- [1] C Green et al, 2012 J. Phys.: Conf. Ser. 396 022020.
- [2] S. Agostinelli, et al, Geant4a simulation toolkit, Nucl. Instrum. Meth. A, Volume 506, Issue 3, Pages 250-303, 2010.
- [3] C. Andreopoulos et al, The GENIE Neutrino Monte Carlo Generator, Nucl. Instrum. Meth. A, Volume 614, Pages 87-104, 2010.
- [4] J. Caballero et al, "OASIS: a data and software distribution service for Open Science Grid" to appear in CHEP 2013.
- [5] Jakob Blomer, Predrag Buncic, and Thomas Fuhrmann. CernVM-FS: delivering scientific software to globally distributed computing resources. In Proceedings of the first international workshop on Network-aware data management (NDM '11). ACM, New York, NY, USA, 49-56, 2011.