Report Title: ERCAP Requests Details
Run Date and Time: 2020-10-05 13:11:32 Pacific Daylight Time
Run by: PDF Generator User
Table name: u_ercap_requests

ERCAP Requests

This request is a renewal.:	true
ERCAP Number of Request to Renew:	ERCAP0014192
ERCAP Number:	ERCAP0017981
Project Title:	Enabling HEP Intensity Frontier Science through HEPCloud
Label:	FNAL-IF
PI Name:	Norman, Andrew (ajnorman)
PI Name Company:	Fermilab
PI Name Email:	anorman@fnal.gov
PI Name Business phone:	630-840-4016

Allocation Year:	2021
State:	Draft
Revisions required:	
Rejection Reason:	
Project Class:	DOE Mission Science
Program:	HEP - High Energy Physics
Sub Program:	HEP - High Energy Physics
Science Category:	Physics : High Energy Physics (Experment)
Project:	m3249

Personnel

Senior Investigators:

Andrew Norman, Fermilab, (anorman@fnal.gov)
James Amundson, Fermilab, (amundson@fnal.gov)

Authorized Preparers:

Norman, Andrew (anorman), Herner, Kenneth (kherner), Holzman, Burt (burt), Kirby, Michael (kirby)

Funding

DOE Office of Science (SC):	true
Funding Office:	DOE Office of Science High Energy Physics (HEP)
Funding Program Manager:	Chatterjee, Lali (lali)
DOE/SC Grant, PAMS, FWP numbers:	DE-AC02-07CH11359
Federal Agency other than DOE/SC:	false
Other Funding/Agency Source(s):	
Other Federal Agency grant numbers:	
LDRD Funding:	false
Funding Laboratory:	
LDRD grant numbers:	
State or local government or agency:	false
State/Local Govt/Agency name and grant numbers:	
Foreign Government or Agency:	false
Foreign Government/Agency name and grant number:	
University:	false
University name and grant numbers:	
Non-profit Organization:	false
Non-profit Organization and grant number:	
Other:	false
Other organizations and grant numbers:	
Office of Science relevance:	

Security

I attest that this project adheres to these guidelines. :	true	I request an exception to these policies, based on the following::	false
Please Explain Policy Exception:			

Project Details

Project Summary and Goals:

The HEPCloud Facility is an intelligent front end interface that is used by High Energy Physics to provision a broad class of resources spanning HPC centers, commercial and academic clouds, and commodity batch clusters. It does this through a dynamic decision engine that is able to use knowledge of allocations on HPC facilities, cost metrics on commercial clouds, and other heristics to make on-demand provisioning requests. HEPCloud also provides a common and transparent interface to the physics experiments using it which enable them to match the requirements of their scientific workflows (CPU core counts, memory requirements, accessability of GPU accelerators etc...) to compute resources at different sites. In particular HEPCloud has provide the interface for the DUNE, NOVA, Muon g-2, and Mu2e experiments to use NERSC's machines (currently Cori and Cori-KNL previously Edison) to perform critical computations.

HEPCloud in 2021 will continue to be the interface that is used to provide resources to neutrino science and precision muon science experiments, which are collectively referred to as the "Intensity Frontier". The major experiments of the Intensity frontier who are using HEPCloud are described below:

The DUNE experiment is future flagship of the DoE Neutrino Program. It is composed of two separate detector complexes - the near detector situated at Fermilab, and the far detector at the Sanford Underground Research Facility in Lead, SD. The experiment is in the process of final design of the near detector complex and its associated detectors. The design of these detectors directly impacts the total sensitivities of the oscillation physics program. The computation of the experiment's sensitivities is required for inclusion in the DUNE Near Detector Technical Design Report (TDR). This is scheduled for completion in 2021.

The NOvA experiment is the currently operating flagship of the long-baseline neutrino science program. The experiment utilizes an enormous 14 kiloton far detector to measure the properties of neutrino interactions, and to observe the appearance and disappearance rates of neutrino and anti-neutrino species as they propagate through the Earth's crust from their origination at Fermilab to Ash River Minnesota, 810 km away. The collaboration analyzes this data to extract information on the fundamental properties of neutrinos and on the model of neutrino oscillations. NOvA has published the world's leading measurements on these topics and continues to expand these measurements with new data and new analysis techniques.

The muon g-2 experiment is designed to precisely measure the anomalous magnetic dipole moment of the muon by measuring their precession in a strong magnetc field. A beam of muons with aligned spins is directed into a storage ring that has a very precisely known magnetic field. The magnitude of that precession is directly related to the difference of g from 2, or g-2. The main goal is to test the Standard Model's predictions of this value by measuring the precession rate experimentally to a precision of 0.14 parts per million. If there is an inconsistency, it could indicate the Standard Model is incomplete and in need of revision.

The Mu2e experiment at Fermilab will search for the neutrinoless conversion of a muon to an electron in the field of an atomic nucleus with a sensitivity of approximately 3x10-17, which is about 10,000 times more sensitive than the previous best measurement. The experiment is nearing the end of construction; first commissioning data with cosmic rays is scheduled for March 2021 and first physics data taking is scheduled for late 2022 or early 2023.

All of these experiments require the simulation of large amounts of collision events in order to accurately model physics processes, detector response, and reconstruction algorithms. These simulations include both processes that would arise from new, speculative physics models and those that arise from standard model processes that would be the background to the new physics. In addition, the experiments that already have real data may run centralized processing or user analysis of said data on NERSC resources.

Detailed Description for DOE Managers:

2021 ERCAP Request for High Energy Physics Intensity Frontier Experiments

The intensity frontier program within the high energy physics portfolio of the Department of Energy, Office of Science focuses on the physics of neutrino interactions, neutrino oscillations, precision measurements with muons and probes fundamental symmetries in the lepton sector of the standard model. The DoE intensity frontier program is fronted by two flagship neutrino experiments: DUNE (final planning and initial construction) and NOvA (in continued operations through 2025) and two precision muon experiments: Muon G-2 (in operations) and Mu2e (in construction). Each of these experiments has significant computing needs relating to their physics missions and for which the facilities at NERSC are required to fulfill.

These experiments, along with smaller experiments in the intensity frontier portfolio, have their aggregate computing needs administered and accessed through the HEPCloud computing portal. This portal allows for the resources being accessed at NERSC and at other sites within the DoE HEP computing ecosystem to be optimally utilized by the scientists on the experiments (i.e. it provides intelligent job routing which can direct work to NERSC when the scale or unique capabilities of the facility are needed, but to other sites when there is spare capacity on them or when the workflows do not require the NERSC capabilities) As a result the ERCAP request for 2021 is being made and aggregated through the HEPCloud project on behalf of the participating experiments. The requests for 2021 are detailed and broken out by experiment as follows:
NOvA (neutrino physics, operational flagship, host lab: FNAL)

The NOvA experiment will continue its scheduled data-taking in 2020 and 2021 with an emphasis on building on the publications and presentation of results which were made in 2020 at the International Conference on Neutrino Physics (Neutrino 2020) in June 2020. In 2021 the NOvA collaboration is working with the T2K experiment on a joint analysis which will combine the results from each experiment to provide stronger estimates of the neutrino mixing parameters than either experiment is able to achieve individually. New minimization/fitting techniques and a joint parameter estimation are being developed which are being to the NERSC platforms. In 2020 improved versions of the 3-flavor fits and parameter estimates were performed on Cori (KNL) with a demonstrated 60% improvement in speed and resource utilization compared to the prior 2019 versions of the fits (i.e. improvements made by the SciDAC-4 HEP on HPC project were able to speed up the fits there by reducing the total resource needs on Cori for the analysis) Additional improvements have been implemented in these code bases and will be utilized in the combined fits with T2K. The codes are expected to exhibit a 16x speed improvement over the prior versions. This improvement is required to handle the increased complexity of the joint T2K fits which are expected to increase the computing needs by a factor of 10-12x the needs of the NOvA experiment in 2020 (i.e. the manner in which each experiment performs their parameter estimation, require that the joint estimation used a hybrid method of profiling over the parameter space which is computationally expensive). In addition, NOvA will be preparing a new sterile neutrino result, and a series of doubly differential cross-section measurements in 2021. Each of these analyses has been tuned and demonstrated run on Cori (KNL) and efforts have been made through the NESAP program to prepare the codes to benefit from upcoming Perlmutter hardware. For 2021 NOvA is requesting:

57 million hours requested for 3-flavor full fitting and statistical corrections computation in the combined fits with T2K. This estimate is based on the most recent analysis fits which were run on Cori-KNL in June 2020 using 73.8 million machine hours (38.9 million hours charged). This estimate includes the known 8x speed improvement from the SciDAC-4 program work, and the factor of 12x increase in computation from the inclusion of the T2K data. This analysis is a fully parallelized MPI-based computation and does not run efficiently on traditional grid resources due to optimizations that have been made for HPC platforms, and for NERSC Cori-KNL in particular. This analysis will run in Q3 of 2021, with a most probable target date of mid August.

30 million hours requested for completion and publication of the sterile neutrino search via neutral current interaction analysis. This analysis uses a covariance-based fitting technique that was developed and tuned for execution on Cori-KNL. The fitter has been tested extensively on Cori-KNL and the estimate is based off this testing along with the projected sensitivities of the analysis. This work is expected to occur late in Q1 of 2021.

12 million hours is requested for neutrino cross-section unfolding. The computation of neutrino cross-section unfolding matrices (under multi-universe neutrino flux variation) have been shown to be an ideal application for utilizing the large memory Cori nodes and the high performance of the global filesystem (the application ingests a moderate number of large input files which benefit from being striped across the file system). NOvA has six separate cross section

results which will require this type of analysis during 2021. Each of these unfoldings is estimated to require 1.5-2M hours as has been demonstrated with the nu_mu charged current inclusive result that was presented in August 2020. This unfolding work is expected to occur throughout 2021, where each of the six runs will take 3-4 calendar days to complete.

4 million hours on GPU enabled Perlmutter nodes for machine learning inference evaluation of NOvA the convolutional network classifier when applied to cosmic ray identification and rejection. This application has been tested as part of a NESAP Tier-2 award and is anticipated to be ready to run in Q3 of 2021 or as soon as early access to the Perlmutter system is available.

Total NOvA Request: 103 million hours
(57M 3-flavor joint fit with T2k, 30M neutral current fitting result, 12M neutrino cross section measurements [6], 4M w/ GPU ML training application)

DUNE (neutrino physics, future flagship, host lab: FNAL)

The DUNE experiment is composed of two separate detector complexes - the near detector situated at Fermilab, and the far detector at the Sanford Underground Research Facility in Lead, SD. The experiment is in the process of final design of the near detector complex and its associated detectors. During 2021, DUNE will be focused on finalizing the near detector technical design report, as well as on the analysis of data from the ProtoDUNE detector that ran in a beamline at CERN, and planning optimization for the ProtoDUNE-II detector that is schedule run again at CERN after the Long Shutdown 2 is completed. As a result, for 2020 we are requesting:

For the analysis of ProtoDUNE Run 1 data and generation of supporting simulation for the analysis, DUNE will require 7.5M hours of computation. This request is specifically to analyze the data using an improved suite of reconstruction code which is better able to identify and classify particle interaction in the liquid argon medium. This estimate is based up current performance of the LarSoft and Pandora reconstruction code on Cori and projected to the side of the full ProtoDUNE Run 1 dataset. The performance metrics which have been run on this code show that it benefits from the expanded core counts on Cori-KNL and is already compatible with Nvidia GPU acceleration.

For the near detector technical design report, an additional 2.1M hours of detector simulation are required to perform final optimizations on the designs. This estimate is based upon current simulations being run both at NERSC and on grid facilities at Fermilab.

For the upcoming ProtoDUNE-II runs that are scheduled to be run at CERN, simulation of the new liquid argon detectors (single phase and dual phase designs) is required to prepare, optimize and estimate background rates for the physics runs. These simulation studies are estimated to require 10M hours of computation each for the single phase and for the dual phase detectors (20M hours total) and are based on the current ProtoDUNE/DUNE software performance.

For the improvement of the long baseline fitting software, which is based upon the NOvA CAFAna suite, we are requesting 2M hours of computation which will be used for validation of a set of new fitters and techniques that have been proposed for adoption by DUNE (i.e. they have been developed by NOvA but need to be merged into the DUNE code base and then validated at scale).

In addition DUNE requests a 200 TB disk allocation at NERSC to be used as a managed storage element (using the RUCIO data management software developed by the Atlas collaboration) for staging data in and out of the NERSC facility.

Total DUNE Request: 31.6 million hours, 200 TB disk
Muon G-2 Experiment (muon physics, operational flagship, host lab: FNAL)
The g-2 experiment is in operations at Fermilab and is currently analyzing the data from their 2018/2019/2020 runs. The Muon g-2 experiment requires 15M hours of computation to calculate systematic uncertainty estimates and perform background studies that are crucial for obtaining a result from the experiment at the targeted 140 ppb level. These computations will correspond to the experiment's Run 2 data, taken in 2019 and be used in the physics result from that data taking period.
The resource estimate is based on experience running the g-2 simulation code starting in May 2020 on Cori and on the initial phase of large scale simulations that were run in September 2020. These computations have been projected to the level of statistics that are required to achieve a 140 ppb sensitivity with the 2019 data set.

Total Muon g-2 Request: 15 million hours, Run-2 systemmatics
Mu2e (muon physics, future flagship, host lab: FNAL)

The Mu2e Experiment at Fermilab will search for the neutrinoless conversion of a muon to an electron in the field of an atomic nucleus with a sensitivity of approximately 3x10-17, which is about 10,000 times more sensitive than the previous best measurement. The experiment is nearing the end of construction; first commissioning data with cosmic rays is scheduled for March 2021 and first physics data taking is scheduled for late CY 2022 or early CY 2023.

The Mu2e collaboration requires 3.2M of simulation for final design decision, 2M hours to inform operations planning for final run plan and 11 M hours for computation of alignment, calibration and commissioning of analysis software prior to first data. The estimates are based on the performance of the Mu2e Geant4 based simulation running in full multi-thread mode across 256 cores of Cori KNL nodes (giving a 40x performance improvement over the average performance of the simulation running on Open Science Grid resource in serial mode), and the performance of the code running on the Cori Haswell nodes.

Total mu2e Request: 16.2 million hours

Additional Liquid Argon Time Projection Chamber Experiments
During 2021 we expect to onboard three additional small scale liquid argon time projection chamber experiments to HEPCloud. These experiments are Icarus, SBND, and MicroBooNE. All three experiments are hosted at Fermilab and are performing measurements of neutrino oscillations at short baseline distances.

Each of these experiments uses the LArSoft/art software stack similar to DUNE. Based on our experiences with onboarding these types of simulation code, we estimate that we will need 3M hours of allocation (1M per experiment) to onboard and test at scale the experiment's code bases.

Website URL:

http://hepcloud.fnal.gov
http://www.dunescience.org
http://www-nova.fnal.gov
https://muon-g-2.fnal.gov
https://mu2e.fnal.gov

Accomplishments Summary:

During 2020 this project enabled major results in neutrino physics. In particular in July 2020 new world leading results in 3-flavor neutrino oscillations were presented by the NOvA experiment at the 29th International Conferenceo On Neutrino Physics (Neutrino 2020) (https://conferences.fnal.gov/nu2020/). These results were produced through the final statistical computations utilizing 39 million hours of computation on Cori-KNL. These results are the most precise fits of neutrino/anti-neutrino data to measure neutrino oscillations. This project also produced results for neutrino cross section measurements with NOvA which were similarlly presented at Neutrino 2020.

For the DUNE experiment this project produced critical simulations which were used for the DUNE Technical Design Report and for the published sentitivies of the experiment. The project was also successful in using the NERSC resources for the simulation and analysis of data from the ProtoDUNE experiment which has lead to the first publication of results from ProtoDUNE.

For the Mu2e experiment this project was used to run simulations that were used to optimize the final designs of the experiment and to estimate sources of background to the very rare signal.

Refereed Publications:

B. Abi et al. Long-baseline neutrino oscillation physics potential of the DUNE experiment, e-Print: 2006.16043 [hep-ex] 2020.

Alexandre Sousa et al. Implementation of Feldman-Cousins Corrections and Oscillation Calculations in the HPC Environment for the NOvA Experiment EPJ Web Conf. 214 (2019), 05012

M. Acero et al. NOvA Collaboration, First Measurement of Neutrino Oscillation Parameters using Neutrinos and Antineutrinos by NOvA, Phys.Rev.Lett. 123 (2019) 15, 151803.

Babak Abi et al. DUNE Collaboration, Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV Far Detector Single-phase Technology, JINST 15 (2020) 08, T08010.

B. Abi (Oxford U.) et al. DUNE Collaboration, First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform e-Print: 2007.06722 [physics.ins-det].

Non-refereed materials:

Resources

NERSC Hours Used:	52,768,396	NERSC Hours Requested:	168,600,000
Archival Storage Used (TB):	0	Archival Storage Requested (TB):	0
CFS Storage Used (TB):	0.014	CFS Storage Requested (TB):	200

Justification for Request:

Nova fitting estimate is based the performance of the 2020 fitters and known performance improvements from associated SciDAC-4 project (HEP Data Analytics on HPC) combined with increases in complexity from addition of the T2K dataset and fitting techniques. NOvA neutral current estimate is based upon 2020 performance runs of the new covariance fitter and the 10% data fitting run that was done. This was then scaled to compute that would be required for the full dataset with the addition of new code that has been shown to improve fit stability (but results in a 6x increase in the computation time). The NOvA cross section estimate is based on the actual compute that was used in 2020 for the nu_mu inclusive cross section unfolding and multiplied by the number of cross section results that are scheduled for publication in 2021 which will require the same treatment (4-6 results). The GPU application estimate is based on performance data that was taken from an application that is in the process of being ported as part of a NESAP.

The DUNE signal phase result is estimated based upon our runs of the DUNE code in 2020 on Cori and scaled for to the processing/analysis of the full dataset. The DUNE near detector design runs are based off of preliminary studies that were done on grid computing (with GPU accelerators) and extrapolated to the expected performance on a similar CPU+GPU node. The long baseline fit estimate is based up on what has been required to port the NOvA fitting code, since the DUNE and NOvA experiments share fitting infrastructure. The DUNE ProtoDUNE-2 runs are based off of what was required for the ProtoDUNE-1 experimental runs, but include the addition of the Dual Phase detector.

The Mu2e final design runs are based off of the 2020 design runs as implimented on a Cori-KNL node utilizing 256 of the available cores/threads. The operations planning run is based upon knowledge of the performance of the Mu2e/art code on the Cori platform as is the request for analysis for alignment and calibration calculations.

The g-2 systemmatics run is based off of test runs that were done in september 2020 of the g-2 code and extrapolated to the full 2019 dataset size.

The request for 3M hours for on boarding of new short baseline liquid argon experiments (SBND, Icarus, MicroBooNE) is based upon our experience with onboarding the other experiments with HEPCloud and the need for scale testing. These activities typically require 1M hours computation to fully vet and test the code bases.

Key Events or Deadlines:

The breakdown of our request by calendar quarters (Q1-Q4) is as follows:

NOvA
    NOvA/T2K Fit 57M Q3
Neutral Current 30M Q1/Q2
Cross sections (6) 12M Q1-Q4
ML GPU Applications 4M Q3/Q4

DUNE
    ProtoDUNE-1 Single Phase Result 7.5M Q2
ND Design Studies (TDR) 2.1M Q2
LBL Fits (improvements) 2M Q1-Q4
ProtoDUNE-2 Sim. Single Phase 10M Q1-Q2
ProtoDUNE-2 Sim. Dual Phase 10M Q2-Q3

G-2
    Run-2 (2019 data) Systemmatics 15M Q1-Q2

Mu2e
Simulation for Final Design 3.2M Q1
Operations Planning 2M Q2
Alignment/Calib/Commissioning 11M Q3/Q4

Short Base LAr
    Onboarding (SBN/Icarus/Micro) 3M Ongoing

Need real-time computing?:

false

Explanation for Realtime Computing Needs:

Experimental or Observational project?:

false

Special Requirements:

Codes

Please tell us about your most used/important codes (select up to 5):

4

Code 1 Name:

art

Code 1 URL:

http://art.fnal.gov

Code 1 Description:

Art is an event processing framework for High Energy Physics. The art framework consists of a collection of libraries and a few runtime-configurable main programs, all written in the C++ programming language.

Code 2 Name:

LarSoft

Code 2 URL:

https://larsoft.org/

Code 2 Description:

LarSoft is a toolkit for the simulation and analysis of data from Liquid Argon Time Projection chambers (LArTPCs). The code is in wide use by the neutrino physics community to support the DUNE, MicroBooNE, SBND and ICARUS collaborations and their detectors. The code is written in C++ with GPU optimizations for key transformations and computational kernels.

Code 3 Name:

CAFAna

Code 3 URL:

https://cdcvs.fnal.gov/redmine/projects/novaart/wiki/CAFAna_resources

Code 3 Description:

CAFAna is a framework for fitting and parameter estimation for neutrino oscillation parameters. The code forms the basis of the NOvA and DUNE fitting suites and has been integrated with through DIY (from the SciDAC FastMath Insitute) to provide near perfect strong scaling applications for parameter estimation and statistical confidence interval computations.

Code 4 Name:

DIY

Code 4 URL:

https://www.anl.gov/mcs/diy-doityourself-analysis

Code 4 Description:

DIY is a code platform that is used to provide dynamic scalablity to MPI based applications. The NOvA and DUNE fitting frameworks (CAFAna) heavily leaverage the DIY code suite to provide automatic blockwise decompotion of their computational domains and to perform reduction patterns. The DIY code base is part of the Exascale Initiative and FastMath SciDAC institute.

Code 5 Name:

Code 5 URL:

Code 5 Description:

Supporting Information

Other HPC Support:

Additional Information:

Feedback:

Usage Agreement

Usage Agreement Initials:

AJN

Award Information

Approval State:

Not Yet Requested

Hours Requested:	168,600,000	Hours Awarded:	
Archival Storage Requested:	0	Archival Storage Awarded (TB):	1
Project Storage Requested:	200	Project Storage Awarded (TB):	
DOE PM Notes:			
Approver:			
Award Status:			

Draft

Record History

Computational Allocation Type:		Archive Allocation Type:	
Computational Current Allocation:	70,400,000	Archival Current Quota (TB):	1
Computational Repo ID:	63327	Archival Repo ID:	63329
Computational Resource ID:	1004	Archival Resource ID:	1006
Project Directory Repo:		HPSS Only:	false
Current Project Storage Quota (TB):	22	Renewed by ERCAP Request:	
Project Directory Repo ID:		Program Manager (historical):	
Project Directory Resource ID:		AY Year Start:	January 19, 2021
Project ID:	63322	AY Year End:	January 10, 2022
Sponsoring Organization:	Fermilab	Does PI work at a federal agency or national lab?:	true

Related List Title: Attachment List
Table name: sys_attachment
Query Condition: Table name = u_ercap_requests AND Table sys ID = 9952c4721beb9c10f6a720ebe54bcb03
Sort Order: Created in descending order

1 Attachments

File name	Content type	▼ Created	Created by
ERCAP0017981 Export(2020-10-05 20:01:42).pdf	application/pdf	2020-10-05 13:01:49	anorman