



Project Close-Out Report

Fermilab HEPCloud Facility Project

Version 1.00

5/10/2019

DocDB # 6619

PREPARED BY:

Eileen Berman, Parag Mhashilkar, Tony Tiradani, Krista Majewski

CONCURRENCES:

James Amundson

Date

Project Sponsor

Close-Out Report Revision Log

Revision	Description	Effective Date
V1.0	First update to template	5/10/2019

Table of Contents

- 1. Project Abstract 3
- 2. Project Documentation 3
- 3. Supporting Documentation 3
- 4. Reason for Closing the Project, Phase 3 3
- 5. Project Goals, Objectives and Deliverables 3
- 6. Project Schedule 3
- 7. Project Team 5
- 8. Budget and Financial Information 5
 - 8.1 Materials and Services (M&S) 5
 - 8.2 Staff (SWF) 5
 - 8.3 Total Project Cost 6
- 9. Outstanding Risks and Issues 5
- 10. Operations and Support 7
- 11. Next Steps 8
- 12. Lessons Learned 9

1. Project Abstract

The goal of the Fermilab HEPCloud Facility Project is to extend the current Fermilab Computing Facility to transparently run on disparate resources including commercial and community clouds, grid federations, and HPC centers. The Fermilab HEPCloud Facility will enable experiments to perform the full spectrum of computing tasks, including data-intensive simulation and reconstruction, at production scale irrespective of whether the resources are local, remote, or both. This will also allow Fermilab to provision scientific computing resources in a more efficient and cost-effective way, incorporating “elasticity” that will allow the facility to respond to demand peaks without overprovisioning local resources, using a more cost-effective mix of local and remote resources transparent to facility users.

2. Project Documentation

Most project documentation is available on the project SharePoint site:

<https://fermipoint.fnal.gov/project/fnalhcf>

Project procurement and financial documents are in the limited-access folder:

<https://fermipoint.fnal.gov/project/fnalhcf/Management%20Internal>

Facility security documents are in the limited-access library folder:

<https://fermipoint.fnal.gov/project/fnalhcf/Facility%20Security>

3. Supporting Documentation

Supporting documentation includes papers and reports presented about HEPCloud or a specific aspect of HEPCloud presented in different venues. All papers and reports of this type are listed on the [HEPCloud web pages](#).

4. Reason for Closing the Project, Phase 3

This project succeeded in achieving several goals listed in the charge. Goals achieved also includes obtaining the Authority to Operate (ATO) from the Department of Energy (DOE) and moving the facility into production.

5. Project Goals, Objectives and Deliverables

The Project Charter defined the following goals, objectives and deliverables:

1. The project will collaborate with the commercial cloud, academic cloud, Grid or HPC resource providers it aims to integrate access to.

ACHIEVED: We have continuously worked with the different providers to ensure we use the resources efficiently and securely.

2. The project will update the communication plan.

ACHIEVED: The project communications lead is Marcia Teckenbrock. She updated and reviewed the communication plan and led the execution of its items such as the one-pager and maintaining the hepcloud.fnal.gov website.

3. The project will incorporate financial information and monitoring infrastructure from Phase-I and Phase-II, develop the algorithms and processes necessary to prioritize use of all resources integrated into HEPCloud and extend the Decision Engine implementing this set of algorithms and processes.

PARTIALLY ACHIEVED: The project has developed and delivered a Decision Engine that is integrated with the Fermilab job and resource provisioning infrastructure. Financial information is monitored and available for the decision making process. The Decision Engine needs to be enhanced to take into account the available financial information and use it in the decision making process.

4. The project will extend the HEPCloud Facility to include the necessary contracts, monitoring, processes and procedures to enable elastic use of Google Compute Platform.

ACHIEVED: The project successfully integrated the Google Compute Platform into one of the available resources in the facility.

5. The project will continue to maintain and update the risk registry, continue to provide a top-level clarification of what project goals or deliverables are flexible and what project goals or timelines must be held.

ACHIEVED: The risk register has been maintained throughout phase 3 and re-evaluated the list of goals and deliverables in response to lack of manpower and upcoming project deadlines. You can view the current risk register on sharepoint [here](#).

6. The project will expand its security documentation, policies, procedures and controls to develop and implement additional items necessary to ensure secure operations in a production environment.

ACHIEVED: The project has defined and implemented all necessary security controls, documentation, policies, procedures, monitoring, and auditing to obtain the ATO and transition to production operations.

7. The project will complete the necessary development, documentation, policies, procedures, monitoring and auditing to transition the facility to production operations.

ACHIEVED: The project has developed the decision engine that satisfied the validation requirements. Policies, procedures, and training materials have been

created. Any other project documentation identified by the project has been completed. Monitoring and auditing projects have the functionality needed to transition to operations.

8. The project will complete the development, documentation and processes necessary to obtain Authority to Operate.

ACHIEVED: The project has obtained the Authority to Operate.

9. The project will continue to work with HEPCloud customers to enable their use of the proto-facility for scientific workflows.

ACHIEVED: The project has worked with CMS throughout the R&D phase to insure continued success of CMS workflows using HEPCloud. The project also worked with the NOvA experiment based on the its use case to provision resources in HPC center at NERSC.

6. Project Schedule

The project ran from Jan 2018 through the beginning of March 2019, inclusive.

7. Project Team

The project organization described in detail in the Project Charter v1.0, Section 11, was used throughout the project, with the following exceptions:

- Krista Majewski was added to the Project Leadership Team in Fall 2018, and officially took over Co-Project Management responsibilities from Eileen Berman at the end of January 2019.
- James Amundson took over the project sponsor role from Panagiotis Spentzouris in October of 2018.

8. Budget and Financial Information

We summarize the project costs broken down by M&S (Jan 01, 2018 to Jan 31, 2019) versus SWF (staff costs). Since many organizations contributed funds and effort to make this project a success, their contributions are noted too. SWF financial data is from Jan 01, 2018 through Jan 13, 2019.

8.1 Materials and Services (M&S)

The project consumed cloud computing services from AWS and Google. These costs are summarized in Table 1 below, broken down by topical category. The M&S contributors are summarized in Table 2 below.

Table 1: M&S Costs by Category from 01/01/2018 through 01/31/2019

M&S Cost Categories	M&S Costs
AWS Cloud Services	\$6,746.25
GCE Cloud Services	\$1,205.79
Travel	\$3,098.00
Hardware (8 VMs & 25 bare metals)	\$13,300.00
Project Total M&S	\$24,350.04

The M&S contributors are summarized in Table 2 below.

Table 2: M&S Costs by Contributor from 01/01/2018 through 01/31/2019

M&S Contributors	M&S Costs
FNAL/Scientific Computing Division (SCD)	\$18,883.67
US-CMS Tier-1 on AWS	\$5,466.37
Project Total M&S	\$24,350.04

8.2 Staff (SWF)

The Staff Costs for the project include all of the Fermilab staff who directly contributed to the project, except the effort by experimenters who contributed only to their experiment's adaptation and operation of the HEPCloud prototype services. We consider the latter to be part of the cost of the experiment's doing business in the cloud and on HPC resources, just as they had to adapt to use local computing or Grid computing services in the past. CMS is treated as a direct contributor because they contributed significant staff resources to develop the HEPCloud infrastructure used by all HEPCloud customers. The project staff effort and costs are from January 01, 2018 through January 15, 2019 are summarized in Table 3 below.

Staff resources are measured in FTE-years. By policy we do not explicitly track project staff costs by person to avoid exposing individual salary information. We can measure however the SWF effort for the project, and the staff costs directly in the case of externally contributed funds. We translate between effort (FTE-years) and costs (\$) using average project staff costs with full burdening. Average project staff considered for calculations is \$250,000 per year composing of 52 weeks. The translation introduces some uncertainty in the breakdown of staff costs given the range of salaries of project team members, so

we qualify this to be an estimate of the staff costs, accurate to a few percent. We mark those numbers that are more precisely known in bold. Other values are our best estimates to a few percent.

Table 3: Approximate Project Staff Effort and Costs from 01/01/2018 through 01/13/2019

SWF Contributors	SWF [FTE-yr]	Appr SWF Cost	SWF Notes
FNAL/SCD	1.91	\$477,500.00	FY18
FNAL/SCD	1.14	\$285,000.00	FY19
US-CMS	0.1	\$25,000	FY18: Effort on core project work
US-CMS	0.03	\$7,500.00	FY19: Effort on core project work
DOE/COMPHEP	1.04	\$260,000.00	FY18: Contact - Lali Chatterjee
DOE/COMPHEP	0.28	\$70,000.00	FY19: Contact - Lali Chatterjee
Project Total SWF	4.5	\$1,125,000.00	

8.3 Total Project Cost

Table 4 below summarizes the total project cost estimate based on the Materials and Services (M&S) cost and the estimated staff (SWF) cost. Detailed breakdowns for each of these are in Sections 8.1 and 8.2, respectively.

Table 4: Total Project Cost Summary

Category	Cost
Project Total M&S	\$24,350.04
Project Total SWF (est.)	\$1,125,000.00
Total Project Cost (est.)	\$1,149,350.04

9. Outstanding Risks and Issues

The risks for this project are maintained in the project risk register. They are reviewed quarterly to ensure they are up to date.

The highest-level issue in the eyes of the Project Leadership Team as the project closes are:

- Obtaining the amount and diversity of personnel resources needed to maintain sufficient effort to accomplish goals set forth for Phase 4 of the project. Specifically the continuation of Decision Engine development and support of operation of services in production, including work identified through the onboarding of experiments. Define the interactions with HPC resources: The evaluation of HPC resources is part of future work. This issue includes allocations (cost model), authentication and authorization, and operations interaction.
- We don't have sufficient auditing and monitoring to track usage and costs for individual customers of the Facility. The project as well as the operations and user support groups have continuously improved existing tools and created new ways to monitor usage minimize the risk. This risk should decrease over time.

10. Operations and Support

The end result of Phase 3 was put the HEPCloud facility into production. The operations of the HEPCloud production infrastructure is outlined in the [HEPCloud Roles and Responsibilities RACI matrix](#). The HEPCloud maintenance is on a 8x5 basis, except for computer security and financial alarms defined to be page-able events which are supported on a 24x7 basis.

Onboarding of experiments to use the new production facility is not part of HEPCloud project Phase 3. A series of meetings and discussions have been held with the FIFE and CMS teams to prepare them for the task of identifying experiments to onboard and running the onboarding process.

In the current R&D Operations phase, jobs will be run on the HEPCloud proto-facility only if they (a) add significantly to our knowledge of the facility and infrastructure and (b) are funded. Requirements for experiments to request to run jobs on HEPCloud include:

- Motivation for utilizing commercial or HPC resources for the experiment and for the project
- Specifications of the campaign, including any items relevant to cost
- Commitment of funding (with source) and experiment resources to campaign
- Completion of any pre-arranged grants with cloud or HPC provider

11. Next Steps

The following tasks (in no particular order) will continue to be pursued by the HEPCloud Project Leadership Team for Phase 4 of the project, dependent on effort and resources available.

- Ensure feedback from code reviews can be applied
- Provide decision engine code as open source
- Support for fixing of bugs impacting production operations
- Decision Engine enhancements

- Improve the maintainability of the Decision Engine and the facility as a whole
- Address the issues surrounding storage on cloud and HPC resources
- ALCF
- Initiate a Phase 4 project to continue the HEPCloud program of work in FY19
- ALCF
- Initiate a Phase 4 project to continue the HEPCloud program of work in FY19

12. Lessons Learned

The lessons learned in Phase 3 are located in the [“Lessons Learned Document”](#) on the project SharePoint site.