

The U.S. CMS HL-LHC R&D Strategic Plan

*Oliver Gutsche*¹, *Tulika Bose*², *Margaret Votava*¹, *David Mason*¹, *Andrew Melo*³, *Mia Liu*⁴, *Dirk Hufnagel*¹, *Lindsey Gray*¹, *Mike Hildreth*⁵, *Burt Holzman*¹, *Kevin Lannon*⁵, *Saba Sehrish*¹, *David Sperka*⁶, *James Letts*⁷, *Lothar Bauerdick*¹, and *Kenneth Bloom*⁸

¹Fermi National Accelerator Laboratory

²University of Wisconsin-Madison

³Vanderbilt University

⁴Purdue University

⁵Notre Dame University

⁶Boston University

⁷UC San Diego

⁸University of Nebraska-Lincoln

Abstract. The HL-LHC run is anticipated to start at the end of this decade and will pose a significant challenge for the scale of the HEP software and computing infrastructure. The mission of the U.S. CMS Software & Computing Operations Program is to develop and operate the software and computing resources necessary to process CMS data expeditiously and to enable U.S. physicists to fully participate in the physics of CMS. We have developed a strategic plan to prioritize R&D efforts to reach this goal for the HL-LHC. This plan includes four grand challenges: modernizing physics software and improving algorithms, building infrastructure for exabyte-scale datasets, transforming the scientific data analysis process and transitioning from R&D to operations. We are involved in a variety of R&D projects that fall within these grand challenges. In this talk, we will introduce our four grand challenges and outline the R&D program of the U.S. CMS Software & Computing Operations Program.

1 Introduction

The Compact Muon Solenoid (CMS) [1] experiment at the Large Hadron Collider (LHC) [2] has had a very successful physics program so far with over 1200 scientific papers submitted to date [3]. The success of this physics program has been enabled by the availability of sufficient computing resources to store, process and analyze the data in an efficient fashion. The CMS experiment is designed, built, and operated by a collaboration of close to 200 institutions across more than 50 countries, and comprises roughly 3,000 members, of which close to 2/3 are physicists with authorship privileges on all CMS physics papers [4].

The U.S. makes up about 30% of the authors across a total of ~50 institutions. Both the U.S. Department of Energy (DOE) [5] and the U.S. National Science Foundation (NSF) [6] are supporting research at these universities. The U.S. funding agencies centrally support both the U.S. contributions to the construction of the CMS detector components, and the operation and maintenance of these detector components and U.S. contributions to the software

and computing needs of CMS. The latter is organized within the U.S. CMS Operations Program. The U.S. CMS Software & Computing (S&C) Operations Program [7] is managing the software and computing parts and is referred to as the U.S. CMS Operations Program in the following.

The U.S. CMS Operations Program is responsible for supporting U.S. physicists to effectively participate in the LHC scientific program. Amongst its responsibilities are:

1. Provide, maintain, upgrade, and operate the integrated computational and data infrastructure and software services to support CMS physics data analysis in the U.S., across a distributed high-throughput computing system including the Fermilab facilities and seven university centers.
2. Support key areas of software development, to maintain and upgrade core software systems, workflow and data management services, and support the modernization of physics software.

Because of the importance of software and computing to the overall physics output of the experiment, the U.S. CMS Operations Program has historically made significant investments in the development and operation of software tools and computing facilities. DOE and NSF-supported efforts are central to CMS computing. DOE-supported Fermilab [8] provides CMS's major computing facility outside of CERN. Fermilab has core expertise and provides development efforts for software frameworks and libraries, data management software, resource provisioning software and cybersecurity. NSF support is critical for operating sites at 7 Universities in the U.S. These facilities form the backbone of the sites supporting CMS and host official CMS datasets, both real and simulated, and a significant amount of user data. NSF-funded institutes play major roles in algorithmic software development, workflow management software, data management software and specialized analysis environments. DOE and NSF funds are also used for supporting operations of the U.S. CMS and CMS software and computing infrastructure and HL-LHC R&D activities. The High Luminosity LHC (HL-LHC) promises scientific results of even great impact, currently out of reach of the LHC, thanks to the upgrade in instrumentation, and a much greater beam intensity.

1.1 The Scale of HL-LHC Computing and Estimated CMS Resource Needs

During the 15 year lifetime of the LHC program, software and computing needs have posed a number of significant and increasing challenges, caused by the distributed nature of the integrated computational and data infrastructure, and the steadily increasing data complexity and scale due to continuous advances in LHC instantaneous luminosity and thus increasing particle multiplicities, occupancies, and additional proton-proton interactions per collision. While LHC computing demands and the scale of the computing problem have continuously increased in the past, CMS and the U.S. CMS Operations Program have responded successfully with an ongoing S&C development program to address these challenges.

HL-LHC pushes these boundaries again, in computing scale, compared to today. HL-LHC, which will take data in several running periods from 2029 to 2042 [9], will collect over 10 times the data collected so far at the LHC. Higher intensity proton-proton collisions and new CMS detector components with higher granularity and more channels will increase the challenges for software and computing significantly:

- (1) The number of events to be processed each year is expected to be larger by a factor of three compared to LHC: 150 Billion events per year.

- (2) Data sample sizes will be larger by a factor of five: the projected total disk storage needs of CMS by 2030 approaches 1/2 exabyte.
- (3) Most data are active and need to be held on quasi-randomly accessible storage systems to be processed by hundreds of simultaneous processing pipelines. This requires very large active storage systems for "hot" data, in addition to agile access to "cold" archival storage.
- (4) Large and highly granular databases with temporal dependencies are used for geometry, calibration, run conditions, and other meta-data.
- (5) Physics software in 2023 consists of more than 4 million lines of highly specialized code that features high algorithmic complexity (e.g. manipulation of small matrices for tracking, complex geometry for clustering, etc.) and low computational intensity, thus providing unique challenges to efficiently match code and data throughput to industry-standard accelerators.

CMS has documented [10] its current best estimates of computing needs for the HL-LHC era and the computing model and assumptions made for these resource need predictions. As part of the reviewing process for the HL-LHC detector upgrades, CMS has developed a set of scenarios [11] that should provide the necessary computing resources for the CMS HL-LHC program at a cost similar to that of the computing resource needs for the current program. It also demonstrates that a "build-to-cost" approach can be taken to mitigate risks: in case of less available resources the experiment can partially compensate by increasing physics thresholds e.g. for reconstruction algorithms, and thus initially compromising physics sensitivities, while eventually being able to catch up when additional computing resources should become available e.g. for re-processing data.

These baseline scenarios still carry a variety of uncertainties and risks. Improvements in equipment pricing may not be as favorable as in the model [12]. It may not be possible to reduce the number of promptly reconstructed events as much as planned. Fast simulation may not be sufficient for as much of the physics analyses as we predict, leading to a greater need for processing resources. Not enough analyses may be able to switch to the most compact data format for analysis (NanoAOD [13, 14]), leading to greater resource demands for processing and storage.

To protect against these risks, the U.S. CMS Operations Program has started a focused program of research that aims to meet the computing challenges of the HL-LHC and reduce resource needs further. The goals for these R&D efforts are

- (1) to reach the required level of scalability for computing facilities,
- (2) to modernize software to efficiently use modern computing architectures and accelerators and improve software maintainability,
- (3) to identify opportunities for more disruptive changes in software and data analysis systems to enhance physics outcomes,
- (4) to minimize technology risks, and
- (5) to improve on needs for maintenance and operations efforts.

The LHC Run 3 that started in 2022 provides opportunities for a "dress rehearsal" of many new computing techniques.




Figure 1. Estimated computing resource needs for CMS [10]. Shown are the modeled annual projections of total CPU and disk needs for CMS through Run 4. The estimated needs for each computing model scenario are shown by the blue lines. The gray band shows the projected resource availability for an example scenario that extrapolates the 2021 CMS pledged resources using an annual increase in available resources of between 10% and 20%. This assumes current WLCG cost projections [12] and a warranty + 3 years replacement cycle of hardware.

2 Grand Challenges for HL-LHC Computing and Software

In order to effectively focus and structure the U.S. CMS R&D efforts, we organize the innovation, research, and development needs for HL-LHC computing into the four *Grand Challenges* that encompass the advances needed for HL-LHC computing to succeed:

(1) Modernizing Physics Software and Improving Algorithms

Exploit novel algorithms, including ML/AI, reduce algorithmic complexity, increase computational intensity, and provide core software infrastructure to enable effective use of modern hardware and accelerators. The work is organized in the following work packages:

- Core Software Framework and Software Portability
- Establish Performance Metric and Performance Baseline for Physics Software
- U.S. Contributions to the Charged Particle Tracking Software
- U.S. Contributions to Software for High Granularity Calorimeter
- U.S. Contributions to CMS Advanced Algorithms Work

(2) Building Infrastructure for Exabyte-Scale Datasets

Build infrastructure to archive, store, transfer, and provide access to exabyte-scale datasets. Explore data lakes and custodial storage: establish a technology and cost model for custodial/archival storage facilities which manages operations costs, and optimizes hardware costs. Orchestrate computational services and data access, provide intelligent

network services, and master the integration of national and international computing resources, including High Performance Computing centers (HPCs). Work packages are:

- Produce Exabyte-scale Central Datasets
- Manage Exabyte-scale Network Flows
- Store Exabyte-scale Central Datasets
- Derive Petabyte-size Analysis Datasets
- Provide access to Petabyte-size Analysis Datasets
- Operate Facilities to support Production, Storage and Transfer of Exabyte-scale Datasets

(3) **Transforming the Scientific Data Analysis Process**

Re-imagine how physicists interact with data when extracting science: seize on industry advances in data science, consider required performance parameters like throughput, latency, ease of use, and functionality, and develop the facility and software infrastructure to support thousands of physicists analyzing exabytes of data. Work packages are:

- Provide Column-Wise Analysis Software Infrastructure
- Provide Column-Wise Analysis Facility Infrastructure
- Provide Column-Wise Data Augmentation Infrastructure

(4) **Transition from R&D to Operations**

The R&D program will contribute to several advances in infrastructure particularly in analysis facilities and networking/storage that will need to be rolled into operations. This transition needs to include adequate testing, monitoring, and training/documentation for a successful adoption by the operational staff. Work packages are:

- Develop sufficient Monitoring Capabilities
- Assess Readiness of Systems/Services for HL-LHC for Transition to User Testing and Operation
- Write Documentation for Systems/Services
- Provide Training for Operators and Users

3 U.S. and International Partners

The U.S. CMS Operations Program partners in these endeavors with its host lab, Fermilab, its partner U.S. institutes and a number of organizations and consortia, allowing leverage of expertise and to find synergies between projects.

U.S. CMS has been proactive in developing an ecosystem for computing and software related research and developments, working with U.S. and international partners. The ecosystem includes *research partnerships*, like those with national labs and with CERN [15]; national and international *consortia* like the Open Science Grid (OSG) [16] and the HEP Software Foundation (HSF) [17]; joint and *collaborative projects*, many of which are listed below, and *community efforts* like the joint Blueprint activities with U.S. ATLAS, OSG, DOE Energy Science Network (ESnet) [18], the Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP) [19], and the Snowmass Computational Frontier [20]. This vast and proactively developed ecosystem and our partnerships provide strength, depth, and cohesiveness in the efforts to develop the HL-LHC computing environment for CMS.

We summarize below partnerships and collaborative work areas with external projects.

Within the DOE program the Center for Computational Excellence (HEP-CCE) [21] is an important partner, that brings together the expertise of HEP physicists and applied math and computational scientists from DOE's Advanced Scientific Computing Research (ASCR) [22] program in an alliance of four national labs (Fermilab, BNL [23], ANL [24], LBNL [25]). The three-year project focuses on the current and future exascale HPC systems and enabling HEP experiments to efficiently exploit their capabilities. It supports research in areas of interest that are shared by multiple experiments and leverages the respective strengths of the laboratories. The highest priority for U.S. CMS is to explore techniques that enable HEP software algorithms to be efficiently executed on a variety of advanced computing hardware architectures. These portability techniques will allow the same code to be executed on CPUs, GPUs and other future accelerators, reducing the effort to evolve the code to new and future HPC machines. Another priority is the exploration of advanced I/O capabilities of the HPC systems and how to optimally use them from the experimental frameworks for HEP workflows.

DOE also funded two relevant projects under SciDAC-4 – *HEP Data Analytics on HPC* [26] and *HEP Event Reconstruction with Cutting Edge Computing Architectures* [27], both of which are led by Fermilab. The former is developing tools and techniques to take advantage of HPC resources in analysis. The latter is developing vectorized and GPU-based tracking techniques for use in both the trigger and reconstruction.

The LHC program relies on very high throughput transatlantic networks as part of its international computational and data infrastructure. ESnet is both provider for the networking infrastructure with CERN and other international computing centers, and a partner for the domestic networking and high-throughput virtual circuits with partner universities (LHCONE [28]). Increasingly, the LHC experiments make intense use of HPC centers, and we engage in setting requirements, and plan the network access and integration of HPC centers into the LHC data infrastructure with ESnet. A focus for HL-LHC will be adding networks as another resource that is managed. Apart from processing and storage resources, the goal is to add the management of networking resources to the portfolio of the experiments to optimize science throughput.

At the same level, the NSF-funded IRIS-HEP contributes partnerships in vital R&D activities. IRIS-HEP serves as an active center for software research, functions as an intellectual hub for the larger HEP community-wide software research efforts, and will transform the operational services required to ensure the success of the HL-LHC scientific program. In particular, U.S. CMS is working with IRIS-HEP on the development of innovative algorithms for data reconstruction and triggering; development of data organization, management, and access systems for the community's upcoming Exascale era; and development of highly performant analysis systems that reduce "time-to-insight" and maximize the HL-LHC physics potential.

In addition, IRIS-HEP funds the parts of the OSG that provide critical infrastructure services for the LHC experiments. OSG has partnered with U.S. LHC over many years in deploying novel computing solutions across the distributed high-throughput computing system and ensuring that they can be operated at sufficient scale. This partnership will continue through the HL-LHC era. The HEP Software Foundation (HSF) and the Worldwide LHC Computing Grid (WLCG) [29] have also been important international partners in research and will continue to be so.

4 Summary

U.S. CMS, in partnership with its host lab Fermilab, its seven Tier-2 institutes and the U.S. CMS research and national and international S&C communities, has established a HL-

LHC R&D program to plan for and build the computing services required for full scientific exploitation of the CMS HL-LHC upgrades. The program directly addresses the four Grand Challenges of HL-LHC software and computing: (1) Modernizing Physics Software and Improving Algorithms, (2) Building Infrastructure for Exabyte-Scale Datasets, (3) Transforming the Scientific Data Analysis Process, and (4) Transition from R&D to Operations. It is fully integrated in the CMS plans for HL-LHC and delivers important components to the overall system.

Management and oversight is provided as part of the on-going U.S. CMS operations program. It also contributes to and strongly leverages the wider computing efforts at Fermilab, the U.S. universities participating in U.S. CMS S&C activities, and the Open Science Grid. It is participating in HEP-CCE within the DOE lab complex, and in the NSF-funded IRIS-HEP LHC Software Institute. In addition, it works closely with the HEP Software Foundation and the CERN-based WLCG program.

5 Acknowledgments

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics, and by the National Science Foundation under Cooperative Agreement PHY-2121686.

References

- [1] CMS Collaboration, *The CMS experiment at the CERN LHC*, **2008 JINST 3 S08004** doi:10.1088/1748-0221/3/08/S08004
- [2] L Evans and P Bryant, *LHC Machine* **2008 JINST 3 S08001** doi:10.1088/1748-0221/3/08/S08001
- [3] *CMS Physics Timeline* <https://cms-results.web.cern.ch/cms-results/public-results/publications-vs-time/>
- [4] *CMS Collaboration Statistics* <https://cms.cern/collaboration/people-statistics>
- [5] *U.S. Department of Energy* <https://www.energy.gov>
- [6] *U.S. National Science Foundation* <https://www.nsf.gov>
- [7] *U.S. CMS Software and Computing Operations Program* <https://uscms-software-and-computing.github.io>
- [8] *Fermi National Accelerator Laboratory (FNAL)* <https://www.fnal.gov>
- [9] *Information about the long-range LHC schedule can be found at* <https://lhccommissioning.web.cern.ch/schedule/LHC-long-term.htm>
- [10] CMS collaboration, *Approved HL-LHC resource projections*, [twiki:CMSPublic/CMSOfflineComputingResults](https://twiki.cern.ch/twiki/bin/view/CMSPublic/CMSOfflineComputingResults)
- [11] CMS Offline Software and Computing on behalf of the CMS collaboration, *CMS Phase-2 Computing Model: Update Document*, **CERN CERN-CMS-NOTE-2022-008** <https://cds.cern.ch/record/2815292>
- [12] S Campana et al, *Overview of the WLCG strategy towards HL-LHC computing - April 2020 LHCC review*, **Zenodo (2021)**, doi:10.5281/zenodo.5499655
- [13] K Ehatäht for the CMS collaboration, *NANOAOD: a new compact event data format in CMS* **EPJ Web of Conferences 245, 06002 (2020)** doi:10.1051/epjconf/202024506002
- [14] G Petrucciani et al on behalf of the CMS collaboration, *Mini-AOD: A New Analysis Data Format for CMS* **2015 J. Phys.: Conf. Ser. 664 072052** doi:10.1088/1742-6596/664/7/072052

- [15] *European Laboratory for Particle Physics (CERN)* <https://home.cern>
- [16] *The OSG Consortium* <https://osg-htc.org/about/introduction/>
- [17] *HEP Software Foundation* <https://hepsoftwarefoundation.org>
- [18] *DOE Energy Sciences Network* <https://www.es.net>
- [19] *Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP)*
<https://iris-hep.org>
- [20] *Snowmass 21: Computational Frontier* <https://snowmass21.org/computational/start>
- [21] *High Energy Physics - Center for Computational Excellence (HEP-CCE)*
<https://www.anl.gov/hep-cce>
- [22] *DOE Advanced Scientific Computing Research* <https://www.energy.gov/science/ascr>
- [23] *Brookhaven National Laboratory (BNL)* <https://www.bnl.gov/world/>
- [24] *Argonne National Laboratory (ANL)* <https://www.anl.gov>
- [25] *Lawrence Berkeley National Laboratory (LBNL)* <https://www.lbl.gov>
- [26] *SciDAC-4: HEP Data Analytics on HPC* <https://computing.fnal.gov/hep-on-hpc/>
- [27] *SciDAC-4: HEP Reco on Parallel Architectures* <https://computing.fnal.gov/hepreco-scidac4/>
- [28] *Large Hadron Collider Open Network Environment (LHCONE)*
<https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome>
- [29] *Worldwide LHC Computing Grid (WLCG)* <https://wlcg.web.cern.ch>