IRIS-HEP 2.0 Strategic Plan

IRIS-HEP Executive Board

Circulated Version of the Plan

Plan was circulated in an email

- Anyone who is interested, please let us know
- Forwarding it to other interested people is fine, please do not post the draft publicly
- Ask us for the overleaf document to add comments!
- We will post this to the archive after the next set of updates (~Jan 16th)

Thanks to everyone who already sent in comments!

- And those of you that have tried to get your communities to comment!

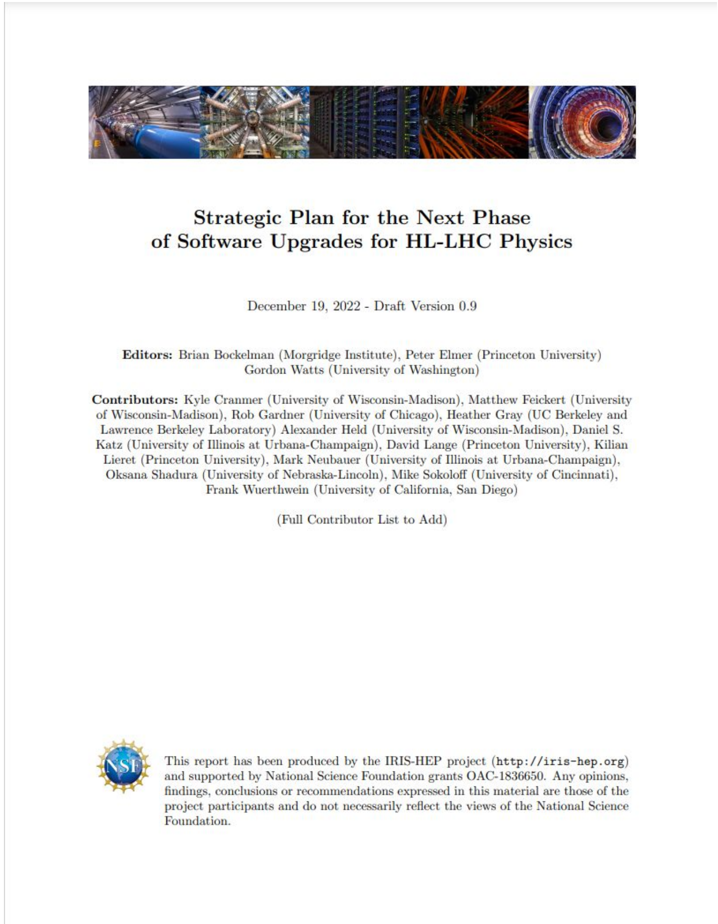

Strategic Plan for the Next Phase of Software Upgrades for HL-LHC Physics

December 19, 2022 - Draft Version 0.9

Editors: Brian Bockelman (Morgridge Institute), Peter Elmer (Princeton University) Gordon Watts (University of Washington)

Contributors: Kyle Cranmer (University of Wisconsin-Madison), Matthew Feickert (University of Wisconsin-Madison), Rob Gardner (University of Chicago), Heather Gray (UC Berkeley and Lawrence Berkeley Laboratory) Alexander Held (University of Wisconsin-Madison), Daniel S. Katz (University of Illinois at Urbana-Champaign), David Lange (Princeton University), Kilian Lieret (Princeton University), Mark Neubauer (University of Illinois at Urbana-Champaign), Oksana Shadura (University of Nebraska-Lincoln), Mike Sokoloff (University of Cincinnati), Frank Wuerthwein (University of California, San Diego)

(Full Contributor List to Add)

This report has been produced by the IRIS-HEP project (http://iris-hep.org) and supported by National Science Foundation grants OAC-1836650. Any opinions, findings, conclusions or recommendations expressed in this material are those of the project participants and do not necessarily reflect the views of the National Science Foundation.

Framing - S&C Gaps

- Do this properly cover the major gaps between now and the start of Run 4, 5+?
- Do this align with experiment priorities?
  - ATLAS, CMS, LHCb
- In one way or the other, we'd like to hang all the activities of IRIS-HEP 2.0 off of these gaps

HL-LHC Software and Computing Gaps

The four software and computing gaps discussed at length in this section are:

G1. **Raw resource gaps**: The HL-LHC dataset will be enormous. Event complexity and count will each go up by about an order of magnitude. If no improvements to algorithms or resource management techniques are made, the HL-LHC experiments will simply be unable to process and store the data necessary for the science program.

G2. **Scalability of the distributed computing cyberinfrastructure**: It is insufficient to buy cores and disk alone – the cyberinfrastructure used by the experiments must also scale to support the volume of hardware. This challenge is especially acute when it comes to data transfers: both the software must be ready and the shared networking resources (e.g., ESNet in the US) must be appropriately managed.

G3. **Analysis at scale**: Analysis at the HL-LHC will be markedly different for two reasons: (a) the scale of the datasets involved and (b) the use of next-generation techniques (such as the latest machine learning techniques) to increase the scientific reach of each result. The former will require users to heavily utilize dedicated 'analysis facilities', optimized for high data rate I/O and the latter will require new services and data management techniques to be developed.

G4. **Sustainability**: HEP is a facilities-driven science - the cyberinfrastructure assembled for an experiment must last or evolve on the decadal scale. This limits some strategies to cyberinfrastructure - for example, it is impossible for LHC to "do it yourself" and own the entire software stack. Specific sustainability strategies must be implemented even at the R&D phase to ensure that the cyberinfrastructure put in place at the beginning of the experiment is one the community can afford.

Strategic Areas of Activity

Analysis Systems

Analysis Systems (AS) builds the tools, libraries, and pipelines that empower a physicist to transform an experiment's production data for physics results.

- Focus:
  - User Experience: AGC as a testbed, adding sys uncertainties
  - Integration of ML into the Analysis Pipeline: an analysis sub-iteration loop (train, tweek, train, analysis, tweek, train)
  - Analysis Reinterpretation: REANA and other tools, deeper integration with new tools
  - Fully Differentiable Analysis Pipeline: tools, interface, study of performance gains

> Providing high levels of interoperability between components of the Analysis Systems ecosystem with each other and with external data analysis libraries is one of the most high-impact activities of the Analysis Systems team.

> The performance, flexibility, and interoperability of the Analysis Systems statistical tools presents opportunities for significant impacts on HL-LHC analysis.

> Complexity of data and analyses at the HL-LHC will grow. Planned improvements to Awkward Array will provide the data handling and manipulation capabilities necessary to approach these challenges.

Analysis Systems - Comments

- OG: Make sure that we tie the skills learned using this ecosystem to people's job prospects in industry
- OG: Better tie in of how we will deal with underlying hardware of analysis facilities and their optimization in this context
- OG: What are the current/planned activities for improving ML integration?
- KB: Is the limiting factor for completing an analysis computing or experiment review processes and procedures
-
- KB: Make sure the tools are in good shape - many people will only use a subset, integration is important, but is not a be-all-end-all

Reconstruction and Trigger Algorithms

- Focus
  - Data Structure Optimization - for use with co-processors *and* CPUs
  - Incorporating Abstract Libraries - low level (like kokkos, matrix calc, etc.) to continue to take advantage of HW improvements over lifetime of HL-LHC
  - Timing & Memory Optimization
  - Exposing Parallelism - for both CPU’s and GPU’s
  - ML for Tracking
  - Integration and Validation - getting the code into the experiment’s production reco and trigger

Amongst the online trigger and offline reconstruction algorithms, the most resource intensive ones are the tracking algorithms: algorithms that together identify where charged particles have traversed the inner part of the detector and their kinematic properties. … Despite being the focus of continuous research during Run 1 and Run 2 of LHC, the compute needs of traditional approaches continue to increase rapidly with increased event complexity and require significant innovation.

Reconstruction and Trigger Algorithms - Comments

- OG: Excellent place to add connections with HEP-CCE
- General: Streamline the number of focus areas
- OG: Missing reference to data structures and optimization
- KB: ACTS can be used everywhere, but is only used by ATLAS, mkfit, etc., is only in CMS - unavoidable?

Translational AI

A dedicated effort in Translational AI will increase adoption of promising new techniques and reduce the time and effort needed to deploy these solutions in the experiments.

Focus:

- Bringing research from the AI community to the HL-LHC - building a self-sustaining pipeline from foundational AI research into the hands of the HL-LHC physicist
- Improving the HL-LHC's interaction with the AI Research Community - Making our data and problems more accessible to the AI research community
- Sustainable AI at the HL-LHC - making sure the AI we do employ will remain well tuned and trained for the life of the HL-LHC

> The faster the community can absorb new AI research and translate it into use by experiments, the more impact it will have on the physics program.

Translational AI - Comments

- OG: This section is fairly abstract and needs more discussion of concrete projects
- OG: Series of blueprints is rather ambitious (too much?)
- OG: Translation fundamental AI work to HL-LHC experiments
- KB: How do we decide what aspects of this fit into the budget envelope? How to prioritize this work vs other aspects of the institute?

Feed these numbers into the current resource requirement modeling and already the disk requirements needed outpace the fixed-budget scenario. The problems compound, however, once one takes into account this data must be moved about the shared national cyberinfrastructure, analyzed at much higher data rates than today at facilities, and done using common infrastructure to reduce sustainability costs

DOMA

Focus

- Scaling up data volumes to HL-LHC - bulk data transfer infrastructure, from benchmarking to replacement and improvement
- Technology Transition for bulk-data movement - tokens, IPv6
- Delivering data to analysis facilities
- Integration Point - DOMA + AS + SSL in a single institute allow us to build out coherent prototypes

> Scaling the cyberinfrastructure to close the 20x gap between today's and HL-LHC's expected transfer rates is one of the most high-impact activities of the DOMA team.

> The prototype Coffea-Casa facility has proven to be a valuable "meeting point" for IRIS-HEP's analysis R&D vision and the team serves as an intellectual hub inside and outside the institute.

DOMA - Comments

- OG: More information on how data is delivered to AF’s
- OG: Where does integrating new columns into existing datasets fit?
- OG: XROOT and its future evolution and development - how do we build a proper community?
-

…activities related to the exploration and innovation of systems, services and physical infrastructure that provide platforms suitable for HL-LHC service environments and runtime ecosystems.

Facilities R&D

Focus:

- Evolving Infrastructure for Next Generation T2 and T3's - The T2 and T3 will transform the historic roles in the WLCG computing hierarchy (to a blended architecture)
- Distributed Infrastructure Management - Managing distributed facilities centrally
- Integration with Production Services (Run 3) - using Run 3 as a test bed

The IRIS-HEP Scalable Systems Laboratory is a Kubernetes DevOps platform for DOMA, Analysis Systems and Innovative Algorithms. The SSL hosts infrastructure for CoDaS-HEP training events and Analysis Grand Challenge pipeline testing and tutorials.

The IRIS-HEP SSL informs the design of next generation WLCG Tier-2 centers and shared Tier-3 analysis facility infrastructure.

Facilities R&D - Comments

- KB: We should be modernizing T2’s and T3’s regardless of the HL-LHC - it isn’t really the driver here - and we should be doing it with or without IRIS-HEP.
- KB: Things like shared T3’s - these are ATLAS concepts - will CMS be interested?
    - Same for metrics section

Fabric of distributed high-throughput computing services

Focus

- Rise of the container-native facilities - “Cloud native” is a generational change, how to pull facilities that are ready along?
- Increasing the diversity of hardware resources - x86, CPU’s, ARM
- Integration with HPC, ML, and non-WLCG resources - No longer dedicated GRID resources only - how best to make them accessible to our field
- Operating Production Services - OSG-LHC

> Established in 2005, the OSG Consortium operates a fabric of distributed High Throughput Computing (dHTC) services in support of the National Science & Engineering community. The research collaborations, campuses, national laboratories, and software providers that form the consortium are unified in their commitment to advance open science via these services.

Training and Workforce Development

People are the key to successful software. People, working together, across disciplines and experiments, over several generations, will be the critical foundation underlying sustainable software.

Focus:

- Basic Curriculum - partnering with HSF, following carpentries model, build out
- Expert-level workshops and mentoring - aimed at people working on reconstruction, etc. Some elements exist, but needs to be made sustainable

Training and Workforce Development

- DK: Introduce career paths when talking about expert level workshops
- DK: Sharper focus on sustainable software development
- KB: Format does not match

Extending several aspects of the software infrastructure will be important for LHCb in Run 4 and critical for Run 5, and may serve as models or starting points for software that is valuable to other experiments during the HL-LHC era.

LHCb/Run 5

Focus:

- GNN's for PV finding - LHCb's operating environment is well suited to the study of this, and some of its algorithms are being ported (e.g. ACTS)
- Modern GPU Hardware - NVidia tensor cores for matrix multiplication - how can we use them better?

Comments

- Format does not match

Grand Challenges

Analysis Grand Challenge

The AGC aims to execute a realistic analysis at the scale and complexity envisioned by the HL-LHC using a set of tools, facilities, and services developed within IRIS-HEP as an exemplar. The AGC team coordinates an annual workshop to demonstrate current progress against the goals and to update the vision and approach as necessary.

- A goal of IRIS-HEP is to help democratize science: anyone, associated with an experiment or not, can re-implement analysis pipelines with the tools and workflows of their choice.
- Systematic Uncertianties, ML Services, Differentiable Analysis, Data Delivery, Data Management, Coffea-Casa, Scaling

> Physics analysis pipelines make use of a large number of tools and services — the Analysis Grand Challenge provides the mandatory end-to-end tests capturing the full complexity of workflows to ensure readiness of the stack for the HL-LHC.

Year	Target
2024	• Define analysis tasks for the top quark mass and di-Higgs measurement. • High-volume analysis done on dataset 20% the scale needed for HL-LHC and completed within 1 hour. • Integrate ML inference service with AGC.
2025	• High-volume analysis done on dataset 40% the scale needed for HL-LHC and completed within 1 hour. • Demonstrate AOD column extraction workflow
2026	• High-volume analysis done on dataset 60% the scale needed for HL-LHC and completed within 1 hour. • Demonstrate fully differentiable analysis
2027	• High-volume analysis done on dataset 80% the scale needed for HL-LHC and completed within 1 hour.
2028	• High-volume analysis done on dataset 100% the scale needed for HL-LHC and completed within 1 hour.

Table 3: Targets by year for the AGC in a five-year timespan.

Data Grand Challenge

The DGC for IRIS-HEP is realized as a set of global data challenges coordinated with the WLCG. These challenges, occurring biennially, bring the entire global community together to demonstrate aggregate transfer data rates and compare what is currently achievable with the HL-LHC roadmap. These challenges also provide an opportunity for integrating new technologies being worked on by DOMA into the production infrastructure.

- The velocity of data is the more vexing issue as the data must be moved over shared network resources (implying coordination and resource management with multiple providers) and the technology stacks select must be interoperable with a global set of facilities.
- Dynamic engineered network paths, tokens, IPv6
- Collaborations: OSG, ESNet

The data challenges are the global community's mechanism to prepare the cyberinfrastructure for the HL-LHC data rates and a mechanism for IRIS-HEP to ready new technologies for the production infrastructure.

Year	Percent of HL-LHC Scale	Minimal Scenario Agg. Targets (Gbps)	Flexible Scenario Agg. Targets (Gbps)
2021	10%	480	960
2023	30%	1,440	2,880
2025	60%	2,880	5,760
2027	100%	4,800	9,600

Data Grand Challenge - Comments

- KB: Has a community as vital as the AGC been built around this?

Training Grand Challenge

To tackle the challenges of the HL-LHC, we need a workforce with broad software knowledge, spanning from basic programming skills to highly specialized training. The TGC defines a roadmap to efficiently scale up training activities and provide adequate training to create the software-skilled workforce that will realize HL-LHC science.

- Democratizing science with standardized prerequisites, scaling up training for intermediate and advanced topics, establishing career paths for developers
- Building a community of educators, sustainable training material,

Year	Target
2018 – 2022	• Develop and teach standardized training modules covering essential software prerequisites (completed). • Establish and grow a training group that coordinates and sustains cross-experiment training efforts (completed).
2023	• Seed topical subgroups that create new intermediate/advanced training material based on the State of Training 2022 survey. • Rebuild and expand the Training Center to become a focal point of *all* software training resources in HEP. • Strengthen collaborations with experiments and new organizations that support the career path of CI Professionals.
2024	The first intermediate-advanced training modules are being taught.
2025	• 80% of all cross-experiment software topics that apply to Ph.D. students should be covered by standardized training modules. • 90% of HEP Ph.D. students should be aware of the material offered by the IRIS-HEP/HSF Training group. • 50% of HEP Ph.D. students should participate in at least one intermediate/advanced training.
2026 - 2028	• Additional focus on workshops, networking opportunities, and discussion platforms for aspiring developers and CI professionals within HEP. • 20% of Ph.D. students and postdocs should be enrolled in software-related communities or have attended advanced workshops. • 20% of HEP Ph.D. students should teach or support a software workshop at one point during their Ph.D.

Training Grand Challenge - Comments

- KB: Has a community as vital as the AGC been built around this?
- KB: How is this a GC as opposed to the regular work that is being done? Are there integrative aspects to this?