Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP)

G. WATTS (UW/SEATTLE)

Physics Goals

Snowmass 2013

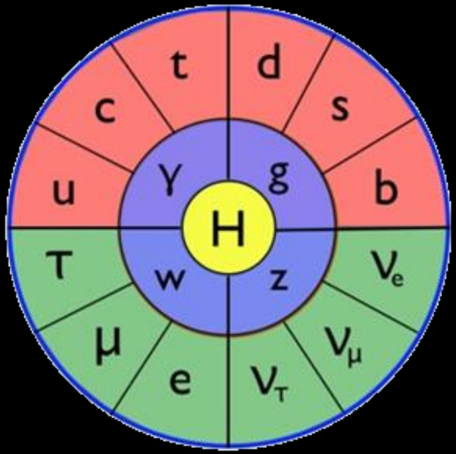

1. **Use the Higgs boson as a new tool for discovery**
2. Pursue the physics associated with neutrino mass
3. **Identify the new physics of dark matter**
4. Understand cosmic acceleration: dark matter and inflation
5. **Explore the unknown: new particles, interactions, and physical principles**

What does the field need to accomplish these goals with the HL-LHC?

Detector & Computing Upgrades!

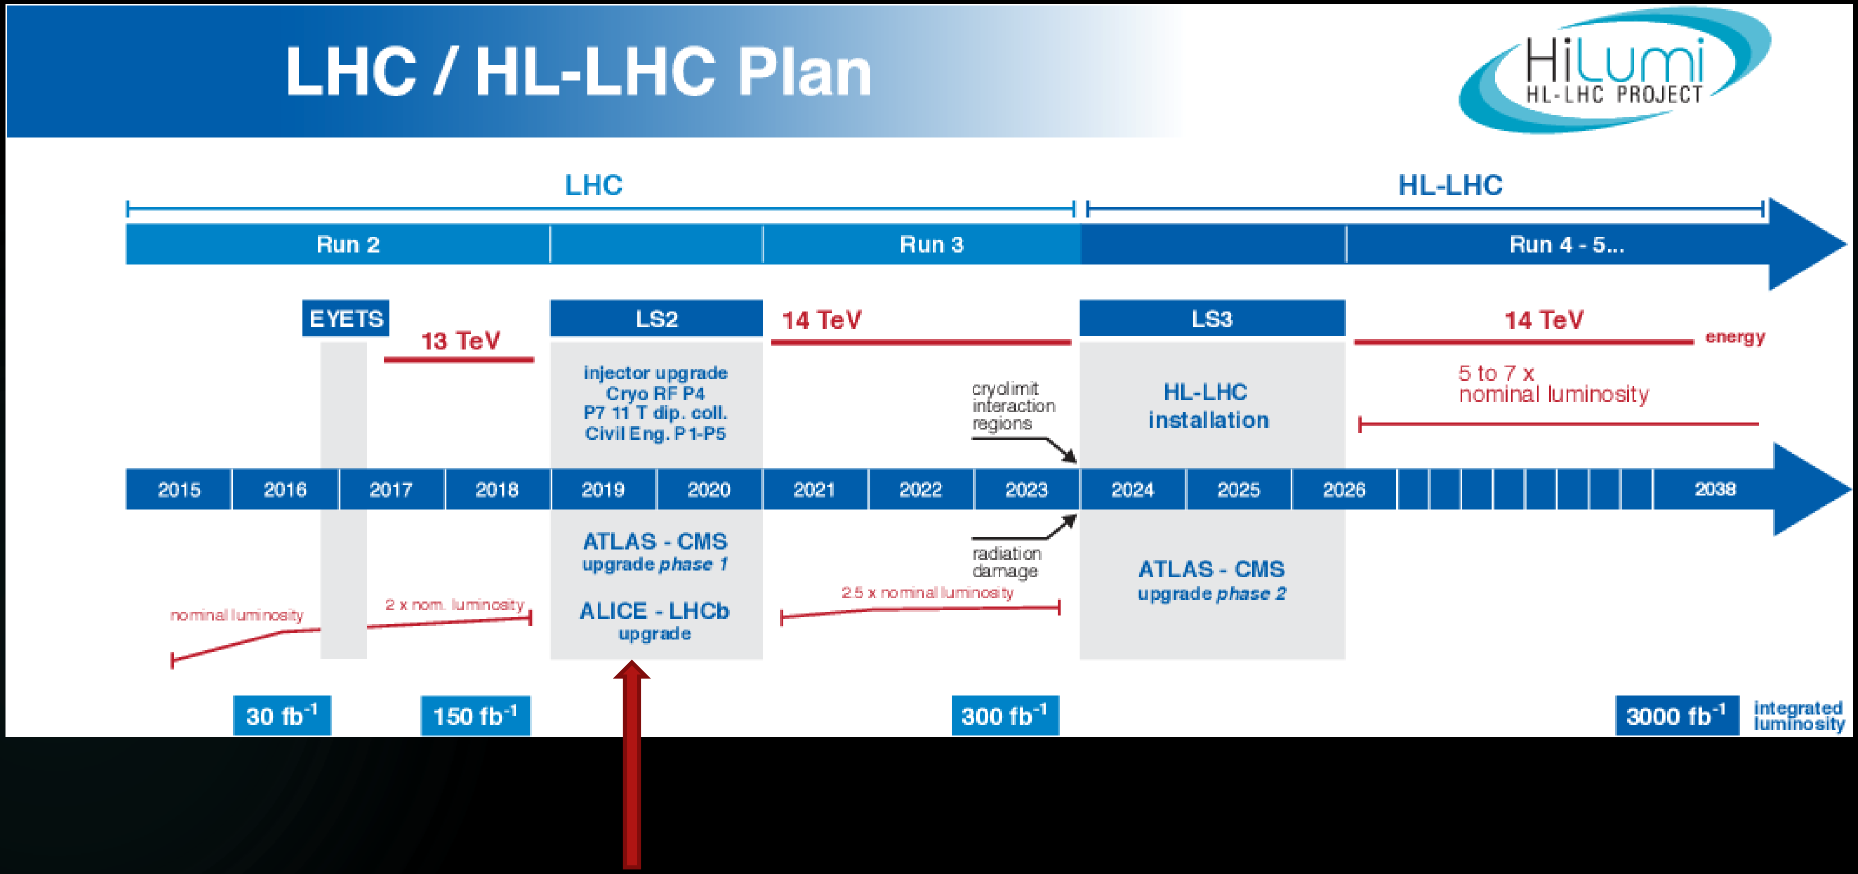

LHC / HL-LHC Plan
HiLumi HL-LHC PROJECT
LHC
HL-LHC
Run 2
Run 3
Run 4 - 5...
EYETS
13 TeV
LS2
14 TeV
LS3
14 TeV
energy
injector upgrade
Cryo RF P4
P7 11 T dip. coll.
Civil Eng. P1-P5
cryolimit interaction regions
HL-LHC installation
5 to 7 x nominal luminosity
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2038
ATLAS - CMS upgrade phase 1
ALICE - LHCb upgrade
radiation damage
ATLAS - CMS upgrade phase 2
nominal luminosity
2 x nom. luminosity
2.5 x nominal luminosity
30 fb-1
150 fb-1
300 fb-1
3000 fb-1
integrated luminosity

Catalyst

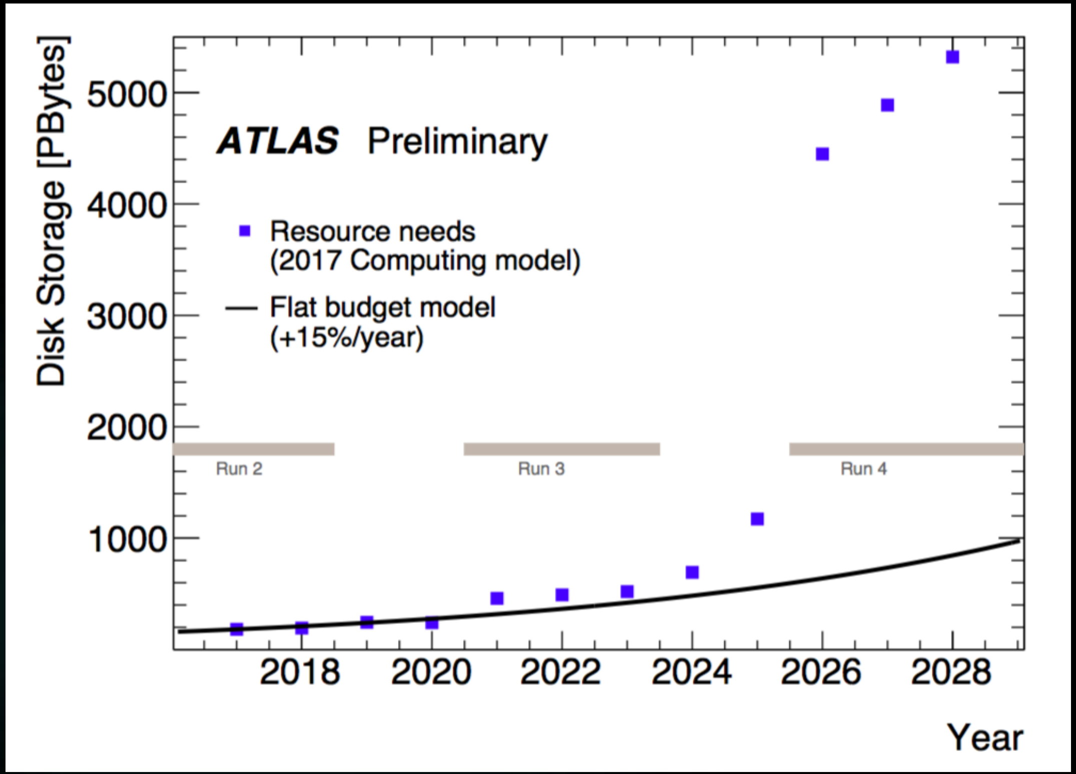
Disk Storage [PBytes]
5000
4000
3000
2000
1000
ATLAS Preliminary
Resource needs
(2017 Computing model)
Flat budget model
(+15%/year)
Run 2
Run 3
Run 4
2018
2020
2022
2024
2026
2028
Year

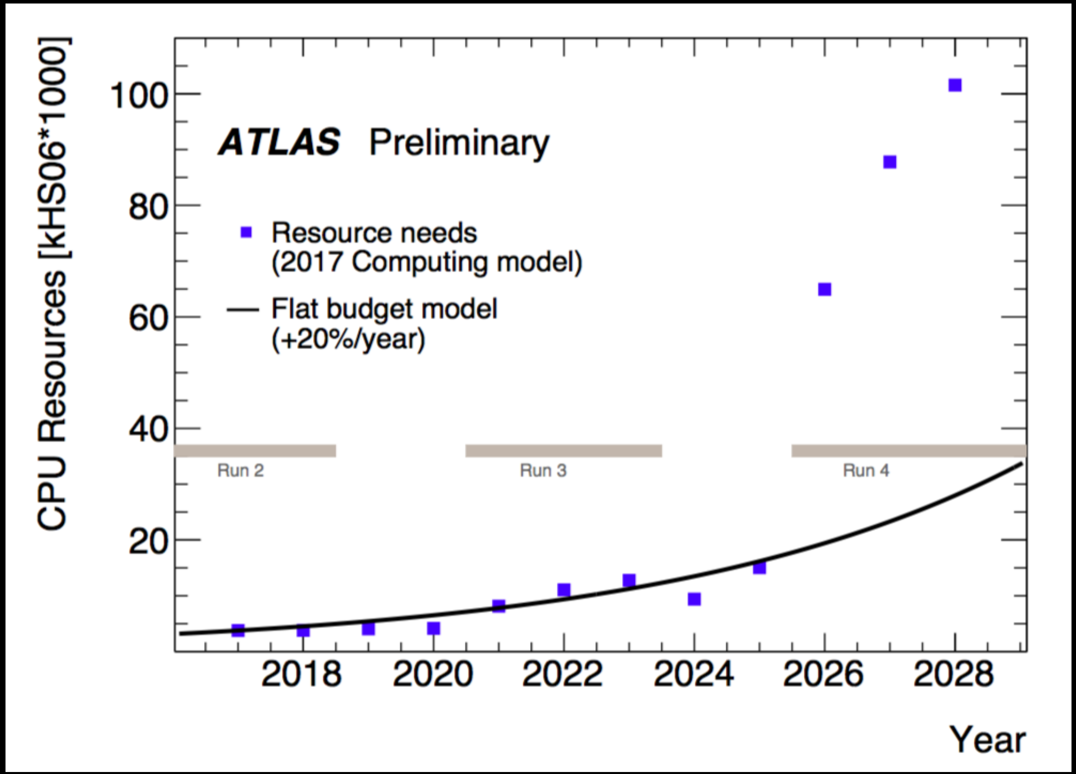
CPU Resources [kHS06*1000]
100
80
60
40
20
ATLAS Preliminary
Resource needs
(2017 Computing model)
Flat budget model
(+20%/year)
Run 2
Run 3
Run 4
2018
2020
2022
2024
2026
2028
Year

Catalyst

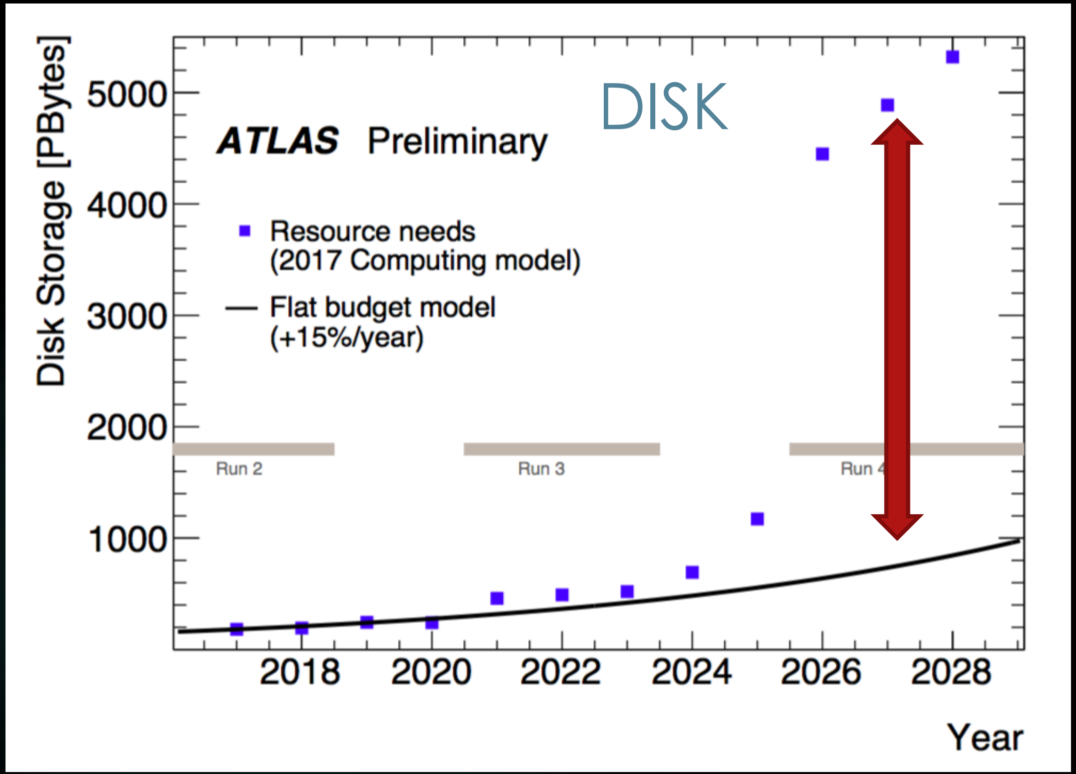

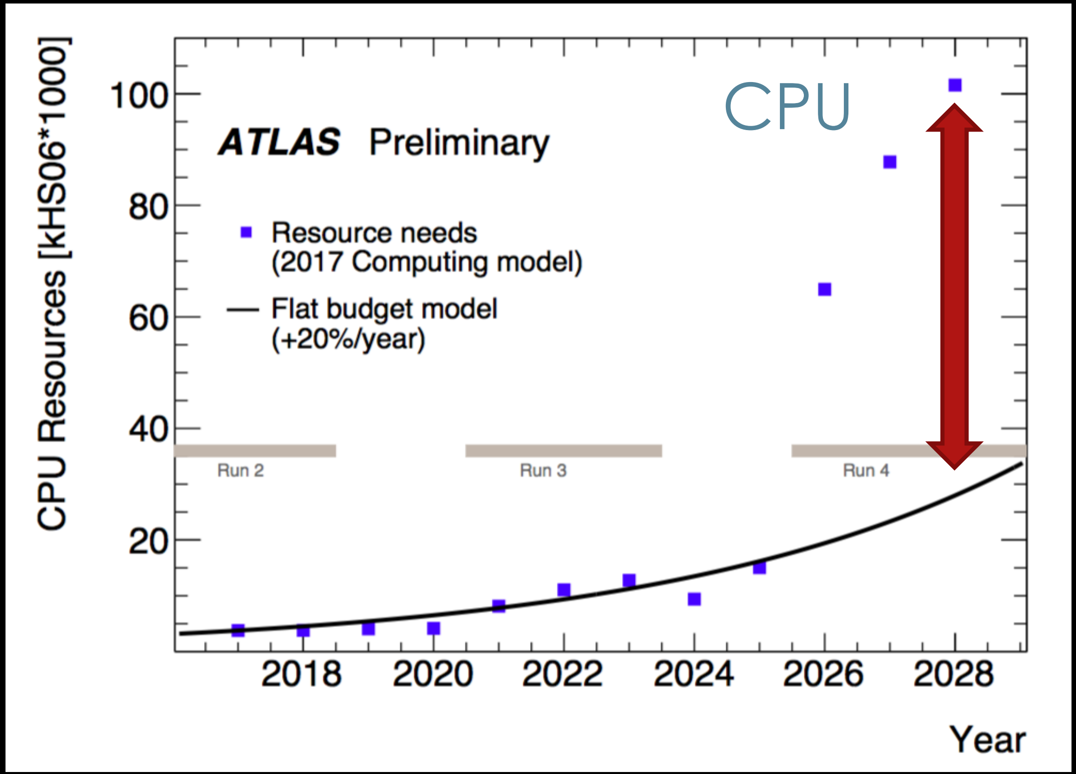

x5 or x10

x5

Both a challenge and an opportunity.

Setting Up The Program

Communities

Building Community Consensus

arXiv 1712.06982

A Roadmap for HEP Software and Computing R&D for the 2020s

HEP Software Foundation[1]

ABSTRACT: Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large investments in detector hardware, either to build new facilities and experiments, or to upgrade existing ones. Similarly, it requires commensurate investment in the R&D of software to acquire, manage, process, and analyse the shear amounts of data to be recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals and priorities, and that the efforts complement each other. In this spirit, this white paper describes the R&D activities required to prepare for this software upgrade.

arXiv 1712.06982

A Roadmap for HEP Software and Computing R&D for the 2020s

HEP Software Foundation[1]

Abstract: Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large investments in detector hardware, either to build new facilities and experiments, or to upgrade existing ones. Similarly, it requires commensurate investment in the R&D of software to acquire, manage, process, and analyse the shear amounts of data to be recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals and priorities, and that the efforts complement each other. In this spirit, this white paper describes the R&D activities required to prepare for this software upgrade.

Contents

Individual Papers on the arXiv

Careers & Training, Conditions Data, DOMA, Data Analysis & Interpretation, Data and Software Preservation, Detector Simulation, Event/Data Processing Frameworks, Facilities and Distributed Computing, Machine Learning, Physics Generators, Security, Software Development, Deployment, Validation, Software Trigger and Event Reconstruction, Visualization

Strategic Plan for a
Scientific Software Innovation Institute (S^2I^2)
for High Energy Physics

Peter Elmer (Princeton University)
Mark Neubauer (University of Illinois at Urbana-Champaign)
Michael D. Sokoloff (University of Cincinnati)

April 6, 2018

- Analysis
- Data Organization, Management, and Access
- Reconstruction & Triggering
- ~Machine Learning

UC-Berkeley, University of Chicago, University of Cincinnati, Cornell University, Indiana University, MIT, U.Michigan-Ann Arbor, U.Nebraska-Lincoln, New York University, Stanford University, UC-Santa Cruz, UC-San Diego, U.Illinois at Urbana-Champaign, U.Puerto Rico-Mayaguez, U.Wisconsin-Madison, U. of Washington, Princeton U, Morgridge Institute iris hep

Institute for Research & Innovation in Software for High Energy Physics

Funded on Sept 1, 2018

http://iris-hep.org

Focus area is where most of the research occurs…

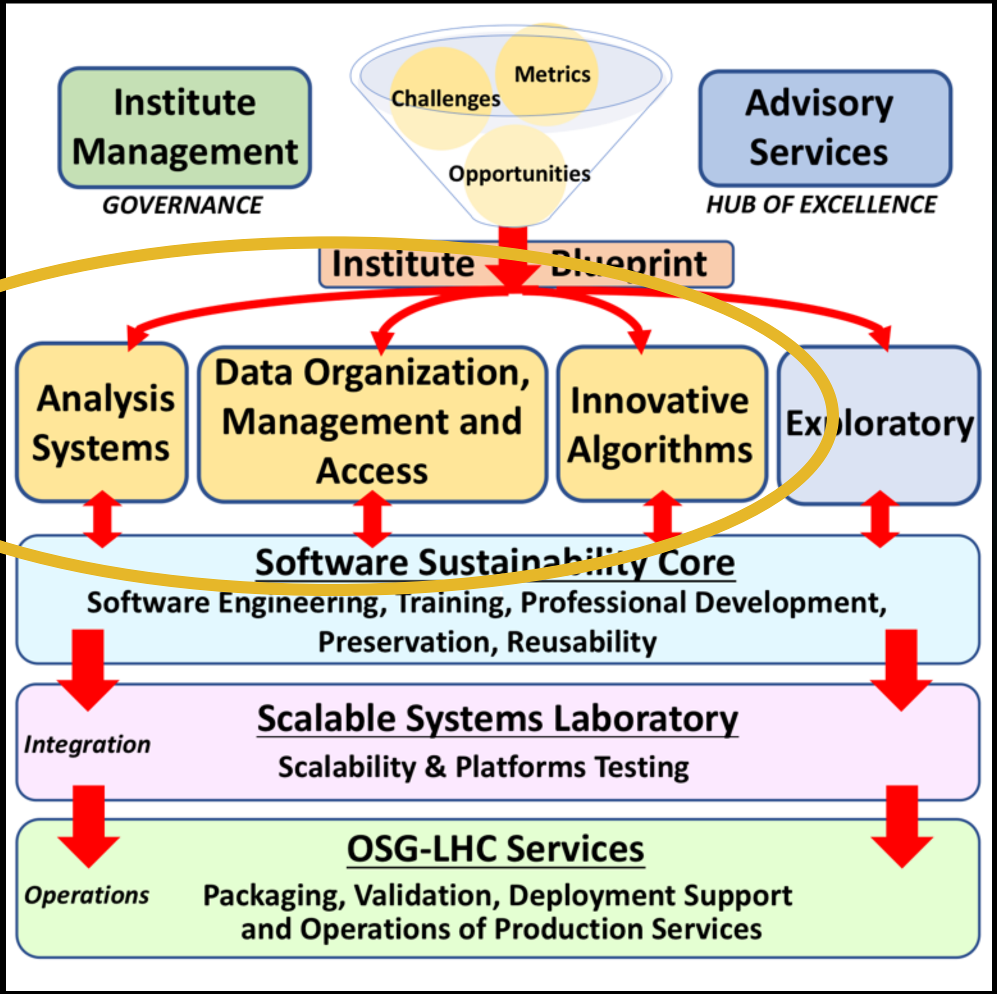

Innovative Algorithms - Trigger/Reconstruction

Algorithms for real-time processing of detector data in the software trigger and offline reconstruction are critical components of HEP's computing challenge.

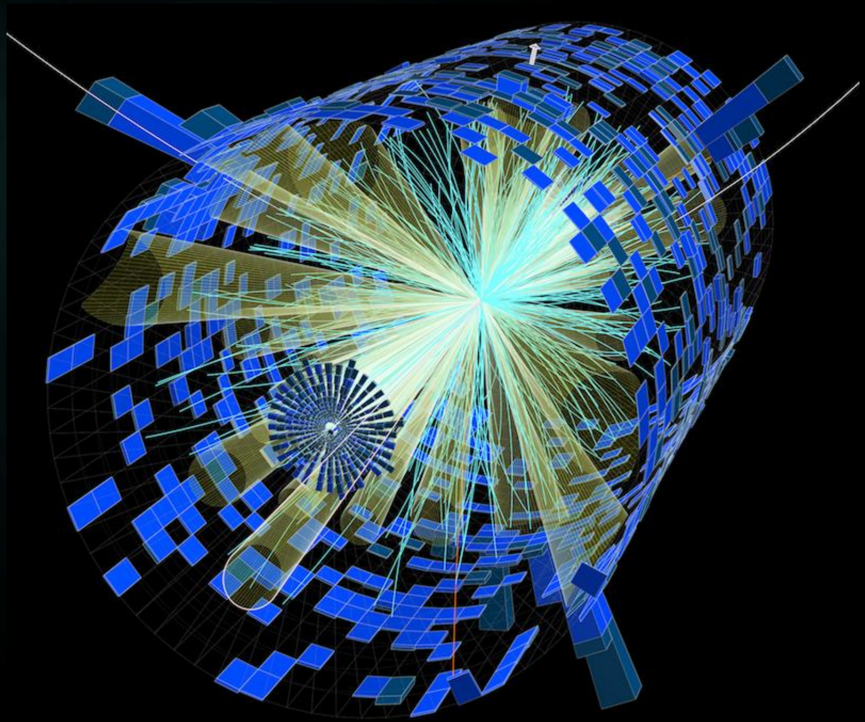

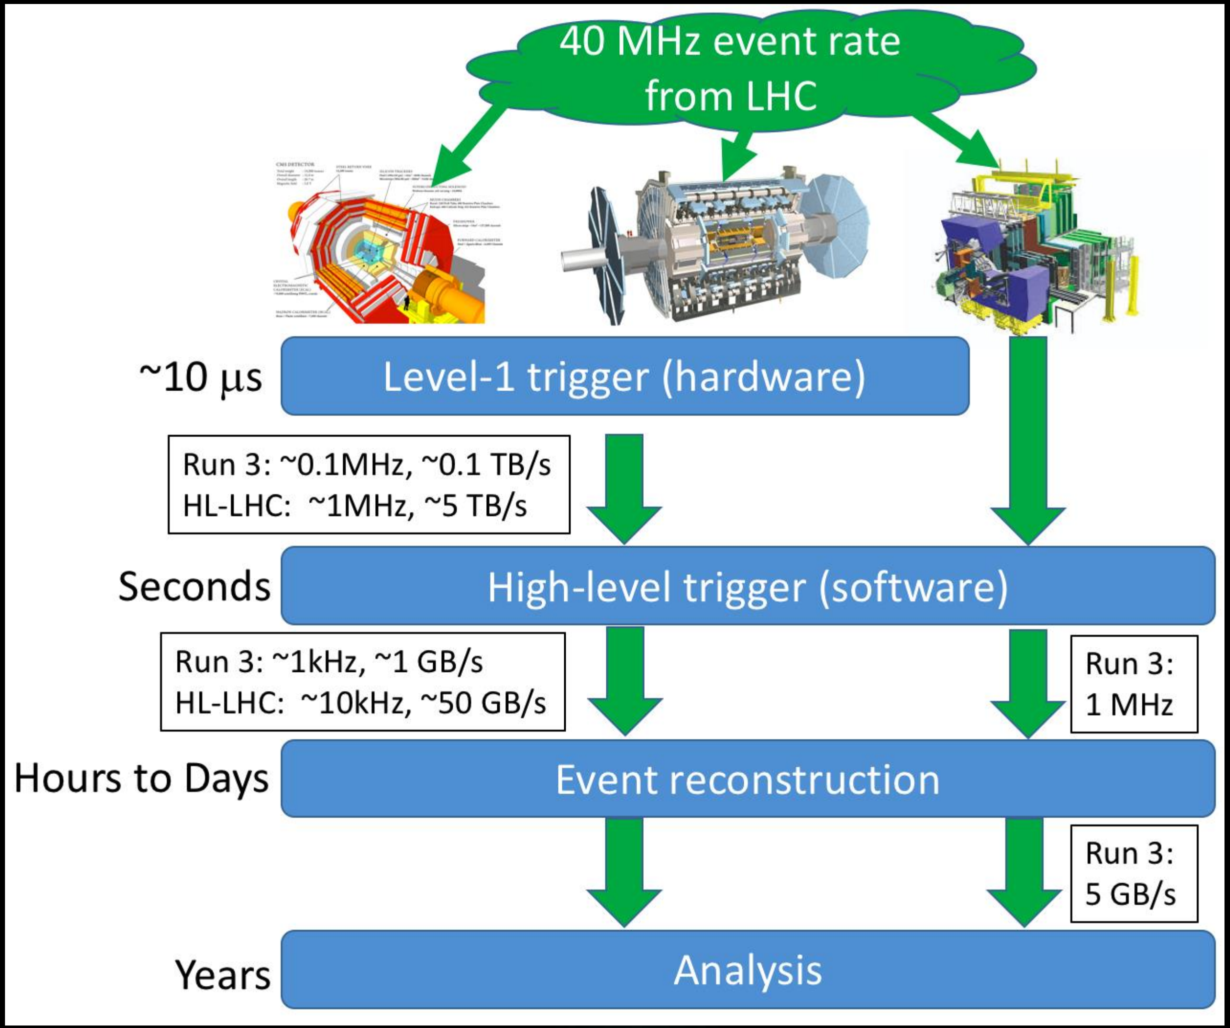

Groups are focused on answering 2 questions

How to redesign tracking algorithms for HL-LHC?

- Determination of charged-particle trajectories ("tracking") is largest component of event reconstruction
- IRIS-HEP investigations
  - More efficient algorithms
  - More performant algorithms
  - Use of hardware accelerators

How to make use of major advances in machine learning (ML)?

- Use of ML in HEP may be a major opportunity
  - Capitalize on industry and data science techniques and tools
  - Could reduce CPU needs
  - Could lead to wider use of accelerators
- IRIS-HEP investigations
  - New HEP applications of ML
  - Use of new ML techniques
  - ML on accelerators in realistic HEP apps

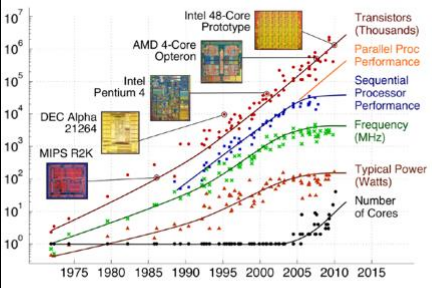

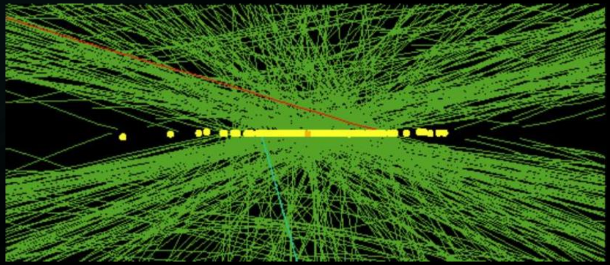

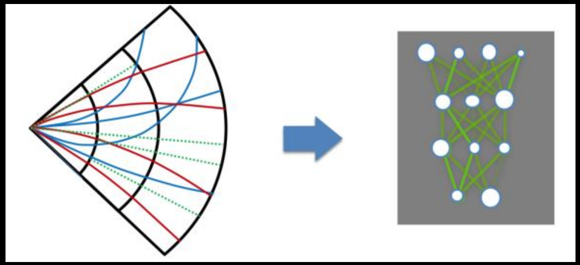

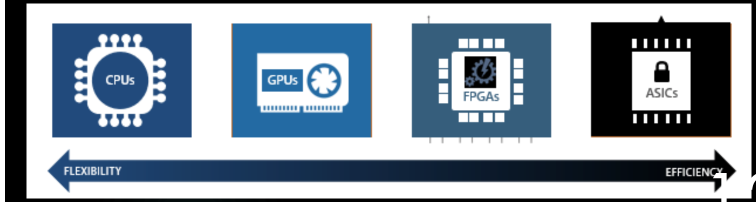

IRIS-HEP Innovative Algorithms Highlights

Parallel tracking contributions to MkFit

- Develop track finding/fitting implementations that work efficiently on many-core architectures (vectorized and parallelized algorithms):
- 4x faster track building w/ similar physics performance in realistic benchmark comparisons

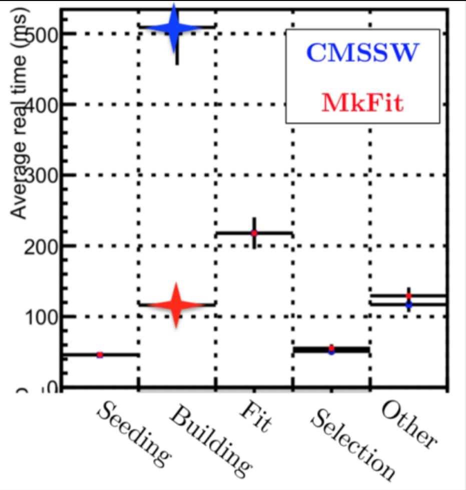

Tracking contributions to ACTS

- Development of the Kalman Filter
- Porting ACTS seeding code to run on GPUs
- Developing connections with other experiments (e.g. Belle-2, JLAB) who may be interested in using ACTS

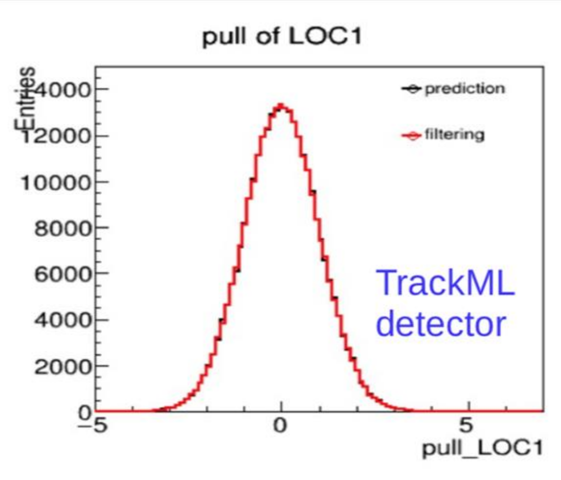

ML on FPGAs contributions to HLS4ML and FastML

identifying specific use cases and operational scenarios for use of FPGA-based algorithms in experiment software trigger, event reconstruction or analysis algorithms

arxiv:1904.08986

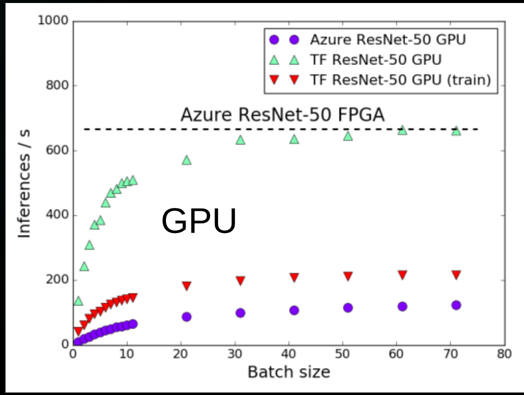

ML4Jets establishing and curating common metrics and data sets

- Aim to connect with diverse segments of machine learning community. Strong connections with theoretical community interested in jet physics
- Tree Neural network approach demonstrated on reference dataset

	AUC	Acc	$1/\epsilon_B$ ($\epsilon_S = 0.3$) single	mean	median	#Param
CNN [16]	0.981	0.930	914±14	995±15	975±18	610k
ResNeXt [30]	0.984	0.936	1122±47	1270±28	1286±31	1.46M
TopoDNN [18]	0.972	0.916	295±5	382± 5	378 ± 8	59k
Multi-body N-subjettiness 6 [24]	0.979	0.922	792±18	798±12	808±13	57k
Multi-body N-subjettiness 8 [24]	0.981	0.929	867±15	918±20	926±18	58k
TreeNiN [43]	0.982	0.933	1025±11	1202±23	1188±24	34k
P-CNN	0.980	0.930	732±24	845±13	834±14	348k
ParticleNet [47]	0.985	0.938	1298±46	1412±45	1393±41	498k
LBN [19]	0.981	0.931	836±17	859±67	966±20	705k
LoLa [22]	0.980	0.929	722±17	768±11	765±11	127k
Energy Flow Polynomials [21]	0.980	0.932	384			1k
Energy Flow Network [23]	0.979	0.927	633±31	729±13	726±11	82k
Particle Flow Network [23]	0.982	0.932	891±18	1063±21	1052±29	82k
GoaT	0.985	0.939	1368±140		1549±208	35k

https://arxiv.org/pdf/1902.09914.pdf

Data Organization, Management and Access (DOMA)

The DOMA focus area performs fundamental R&D related to the central challenges of organizing, managing, and providing access to exabytes of data from processing systems of various kinds.

- **Data Organization**: Improve how HEP data is serialized and stored.
- **Data Access**: Develop capabilities to deliver filtered and transformed event streams to users and analysis systems.
- **Data Management**: Improve and deploy distributed storage infrastructure spanning multiple physical sites. Improve inter-site transfer protocols and authorization

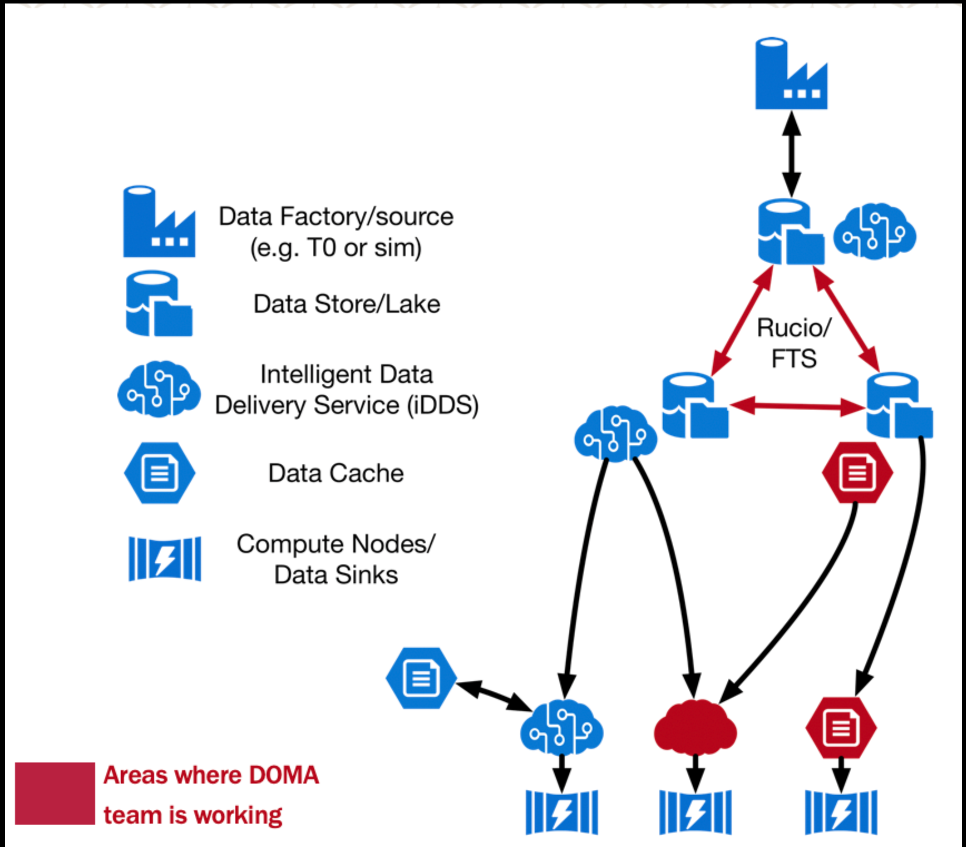

DOMA: Intelligent Data Delivery

- In the HL-LHC era, we must deliver more events - and at lower latencies - if the analysts want to make progress!
  - Low-latency delivery of events requires transformation from long-term archival formats that we want to decrease data size.
  - Data should be transformed and delivered at the storage level, not at the workstation.
  - Users should be enabled to work on a multitude of data formats (esp. non-ROOT) without having to write them to disk.
- We are currently prototyping an Intelligent Data Delivery service to:
  - Extract events from a data lake for fine-grained processing
  - Deliver events to analysis facilities at a high data rate.

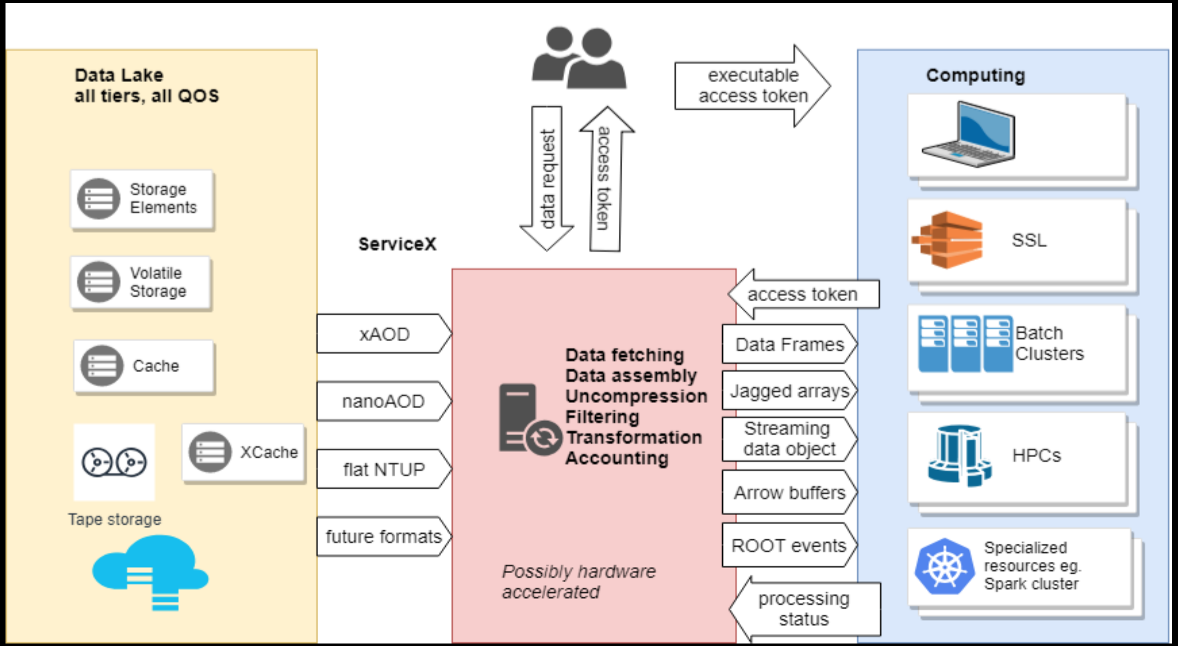

IMPACT / Status:

- Working to integrate intelligent data delivery with ATLAS's PanDA for fine-grained processing.
- Can transform and deliver ATLAS xAOD events for analysis.
- Working with Coffea team to deliver CMS NanoAOD events to Jupyter notebooks.

DOMA: Moving Bulk WLCG Data

- There is a strong movement in the community to move from niche protocols for bulk data movement to more standardized ones such as HTTP.
  - Bockelman co-leads the working group within the WLCG for "third party copy' (TPC).
- During IRIS-HEP, HTTP-TPC has gone from small test transfers to scale tests on servers to scale tests in the WLCG DOMA community.
  - **Demonstrated HTTP's ability to achieve speeds similar to GridFTP on dedicated server hardware.**

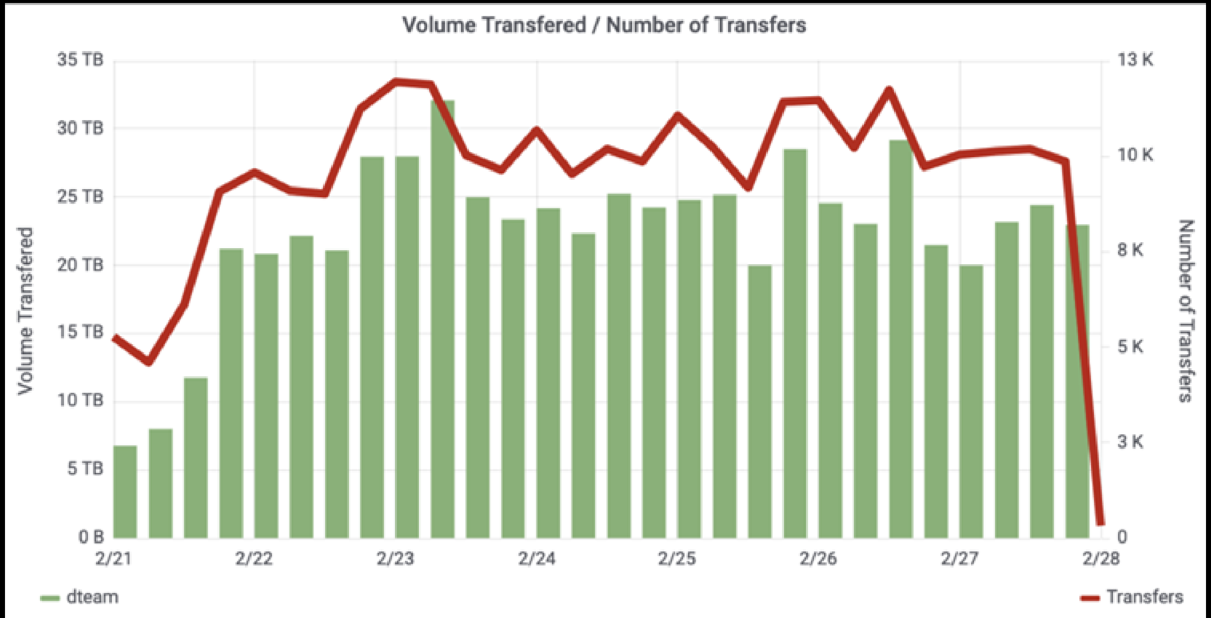

IMPACT / Status:

- Worked to make HTTP-TPC available in the storage systems used by U.S. LHC sites.
- With WLCG, worked to finalize a common, interoperable authorization scheme based on OAuth2 and JWT.

Analysis Systems

Develop sustainable analysis tools to extend the physics reach of the HL-LHC experiments

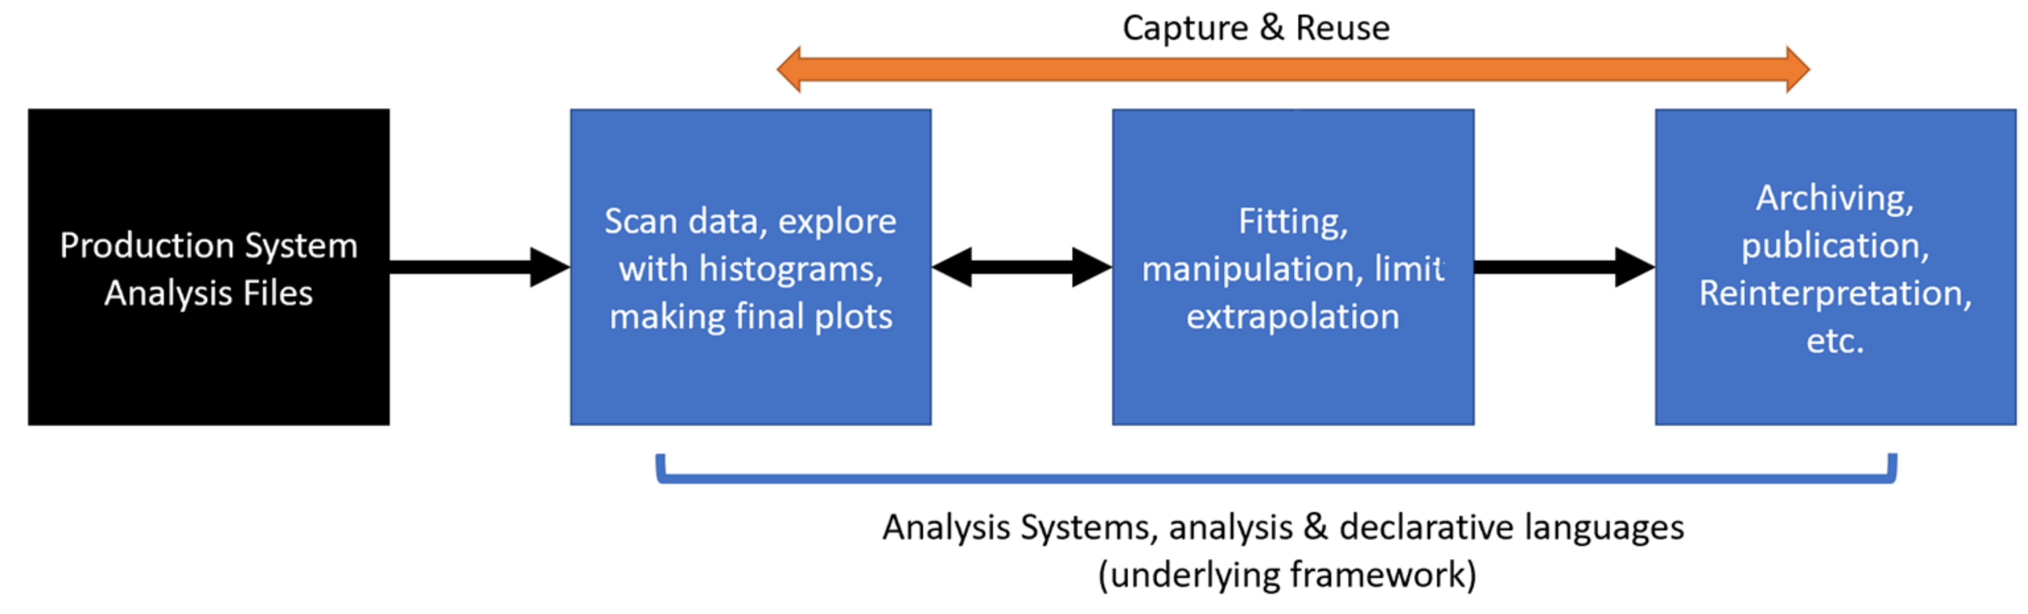

- create greater functionality to enable new techniques,
- reducing time-to-insight and physics,
- lowering the barriers for smaller teams, and
- streamlining analysis preservation, reproducibility, and reuse.

Analysis Systems projects span all stages of end-user analysis.

Analysis Systems - Data Query

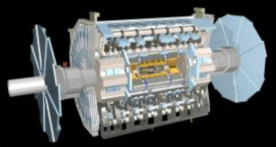

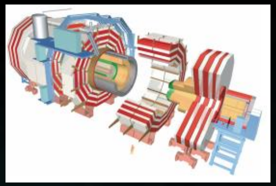

Production & Simulation

IDDS & ServiceX

Designed for high speed data delivery

Awkward Array

Apache Spark

Coffea

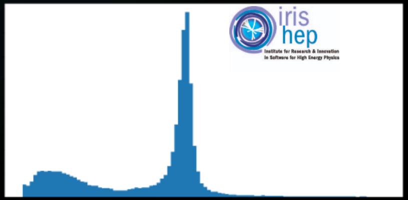

- **ServiceX (IDDS)** - Collaboration with DOMA to serve analysis data in a "columnar form".
- **Awkward Array** - HEP-style column-wise data in the python ecosystem for manipulating the data
- **Coffea** - column-oriented framework for analysis (developed initially at FNAL in the US CMS context)
  - Builds on top of other backends allowing execution on Spark- or HTCondor-based resources.

Full chain to make a Z mass peak in electron data!

Infrastructure

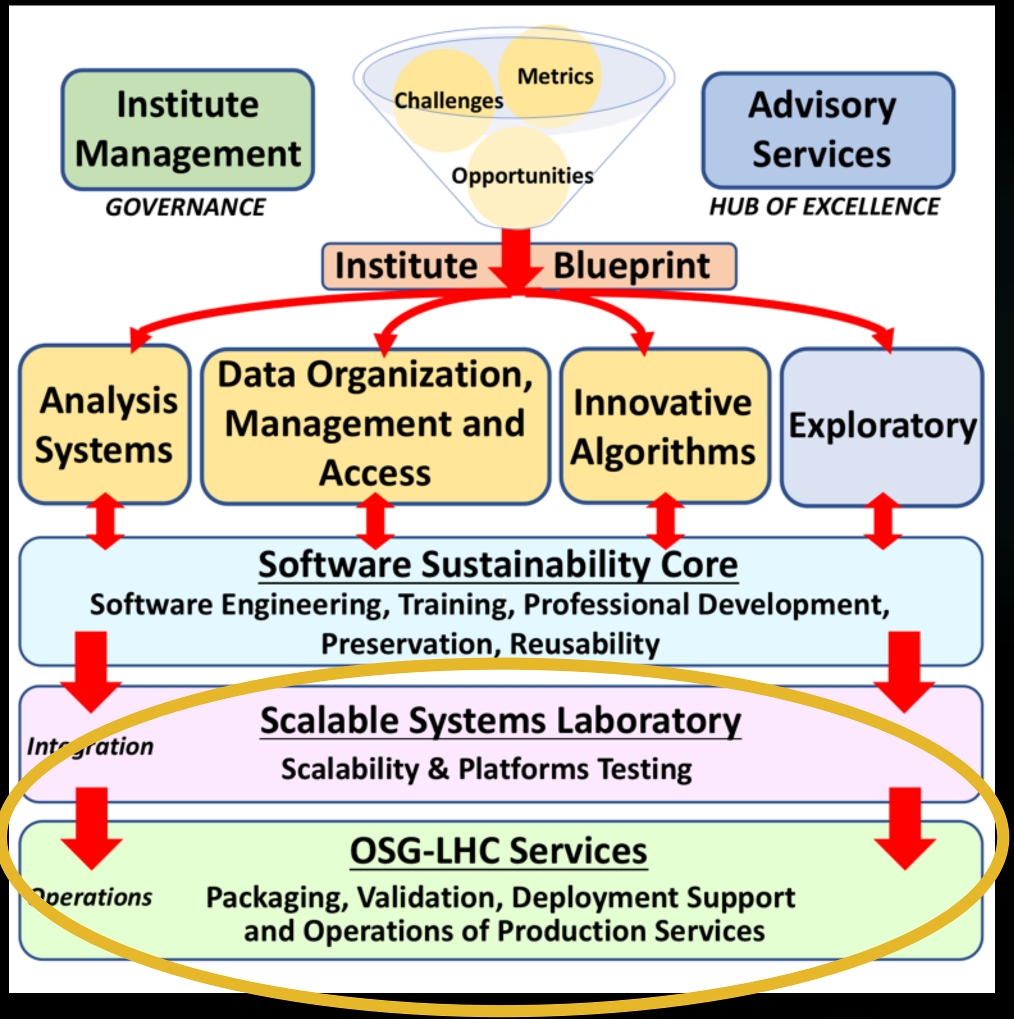

Scalable Systems Laboratory (SSL)

Goal: Provide the Institute and the HL-LHC experiments with scalable platforms needed for development in context

- Provides access to infrastructure and environments
- Organizes software and resources for scalability testing
- Does foundational systems R&D on accelerated services
- Provides the integration path to the OSG-LHC production infrastructure

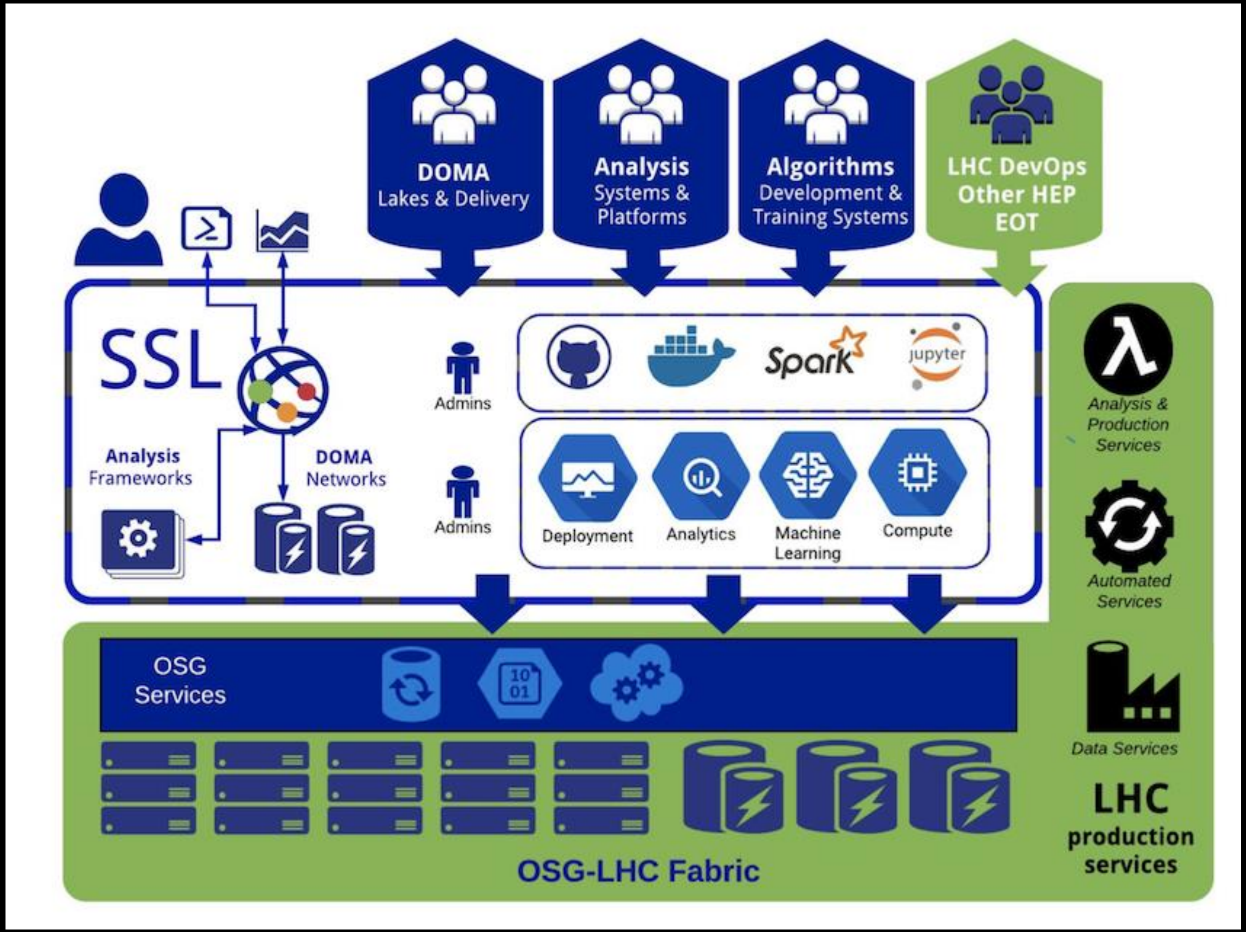

Open Science Grid (OSG)

- The OSG is a consortium dedicated to the advancement of all of open science via the practice of Distributed High Throughput Computing, and the advancement of its state of the art.
- IRIS-HEP is the funding mechanism to support LHC needs from the consortium.
  - Effort form IRIS-HEP is roughly ⅓ of the total effort in OSG today.
- At a high level policy, only shared interests between US ATLAS and US CMS ops programs are within scope of the OSG effort in IRIS-HEP
  - There are activities in the OSG consortium more broadly that serve multiple domains, DOE-NP, and cosmic and intensity frontier experiments in DOE-HEP, plus one of the experiments but not the other. Such activities are not within scope of IRIS-HEP.

OSG Highlight: Transitioning away from Globus

- Globally, the LHC today depends on GSI for authentication and gridFTP for bulk data transfer.
  - Both are no longer supported by their original developers.
- OSG inherited the source code, and is maintaining it within the context of "Grid Community Toolkit" that was created for this purpose in 2018.
- We developed a roadmap for replacement of both GSI and gridFTP that has been socialized globally, and across science domains.
  - August 22nd 2019: Roadmap and schedule presented to LHC ops program via OSG council
  - September 12th 2019: Roadmap and schedule presented to WLCG via GDB
  - January 2020: First demo of a US-LHC site free of GSI and gridFTP (prototype & proof-of-concept)
  - January 2021: OSG production software without GSI and gridFTP is released
  - January 2022: End of support of GSI and gridFTP in OSG releases.

Infrastructure

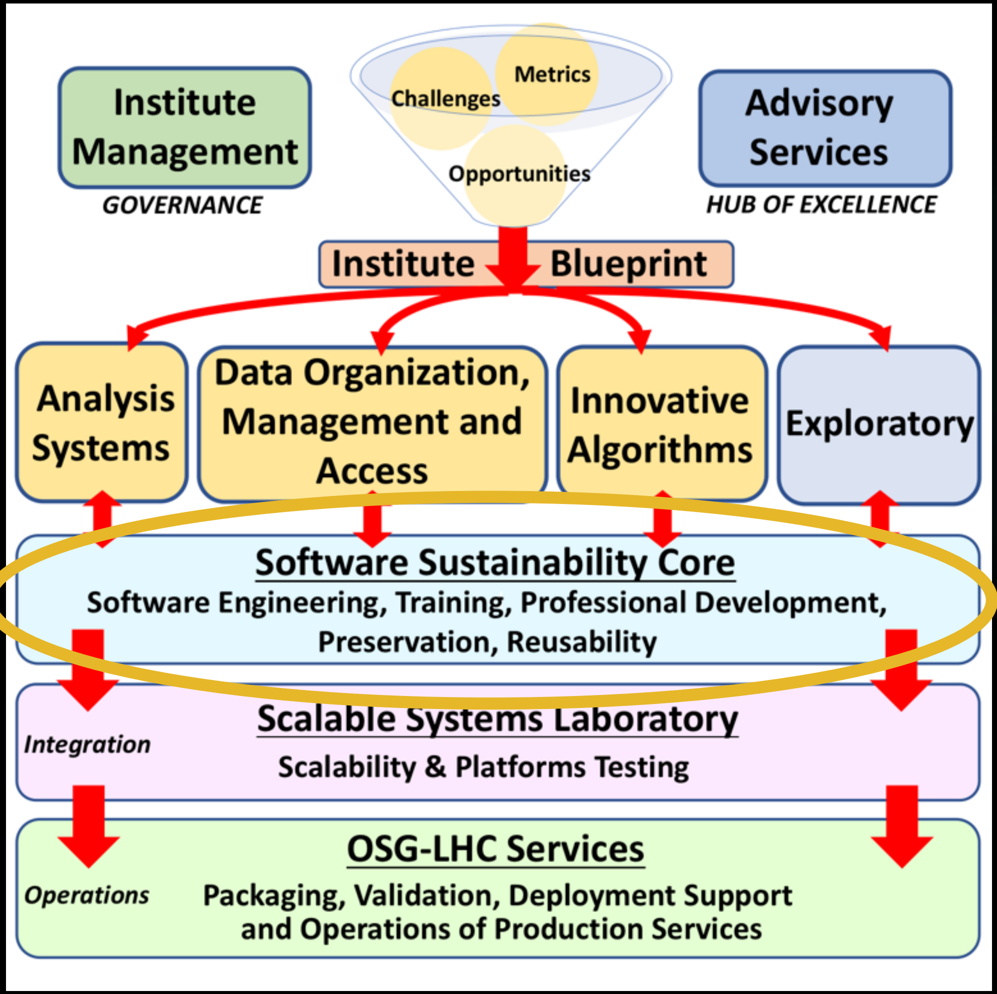

Focus Area: Training

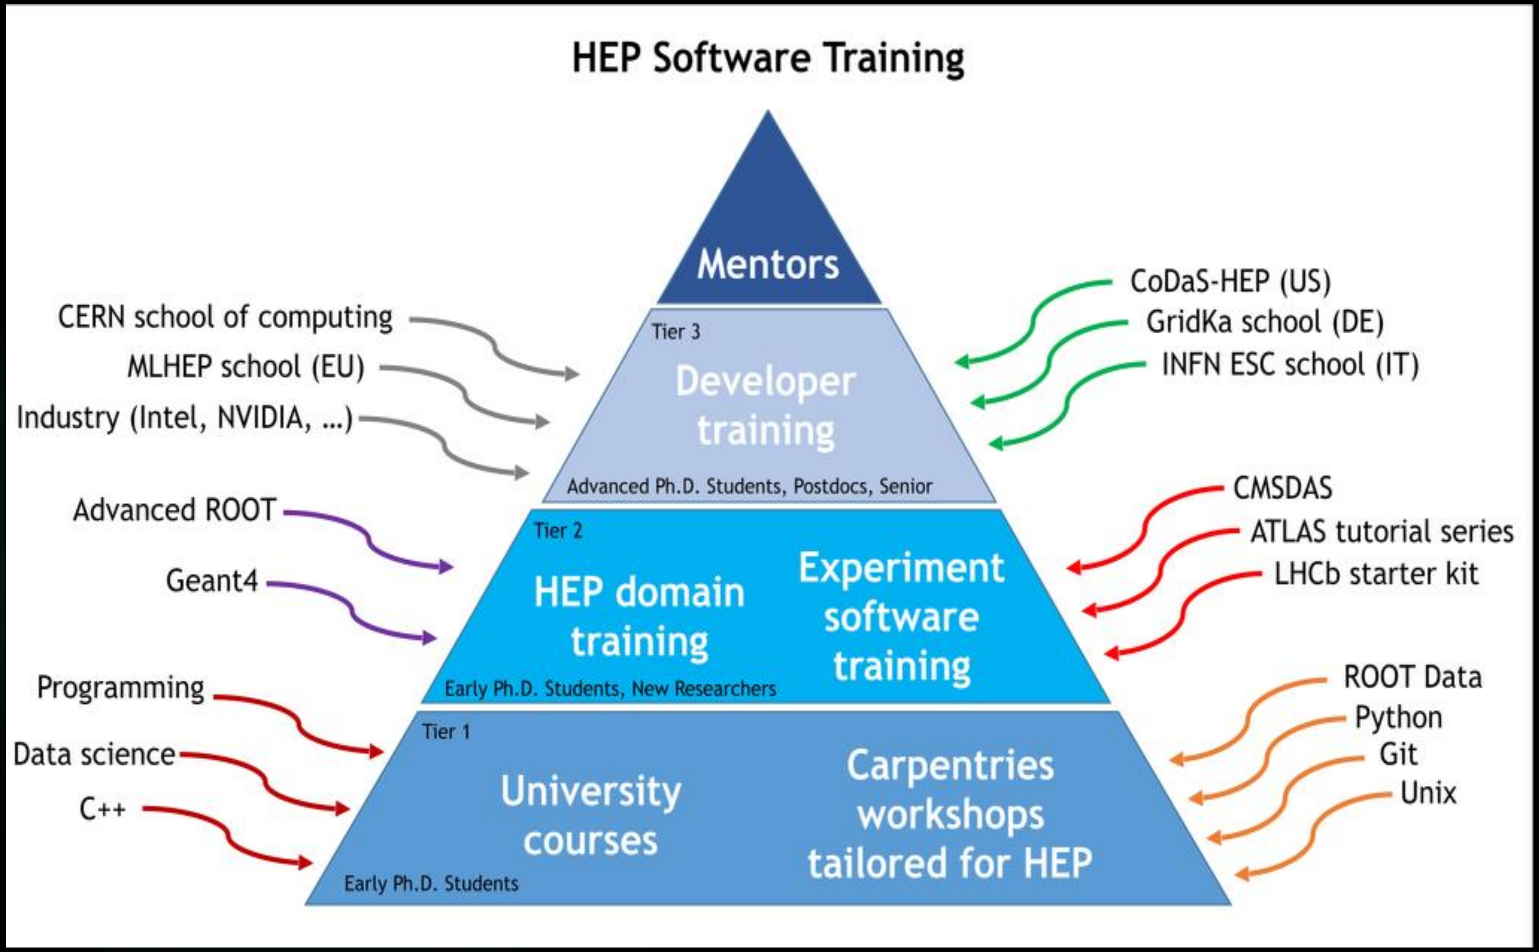

- Our goal is to make training **Sustainable and Scalable**
- This framework recognizes that people need to continue to acquire relevant software skills as their research experience grows.
- The pyramid shows this progression, with expertise building from the bottom up

This is a general framework for training, but from the NSF we have funds from both IRIS-HEP (OAC-1836650) and a separate NSF CyberTraining project FIRST-HEP (OAC-1829707, OAC-1829729, http://first-hep.org) which are working towards implementing this model

Major Activities

- Training for HEP community
  - Software Carpentry Workshops
    - Github/Unix/Python/Plotting
    - Universities, National Labs
    - 180 people trained
    - Female participants highly encouraged
- Outreach
  - Programming for STEM teachers
  - Underrepresented Communities
  - Scientific Software Club at UPRM
- Careers and Jobs
  - IRIS-HEP Fellows (HEP/non-HEP)
    - Graduate
    - Undergraduates
- Collaboration
  - FIRST-HEP (*http://first-hep.org*)
  - HEP Software Foundation (*https://hepsoftwarefoundation.org/*)
  - The Carpentries (*https://carpentries.org*)

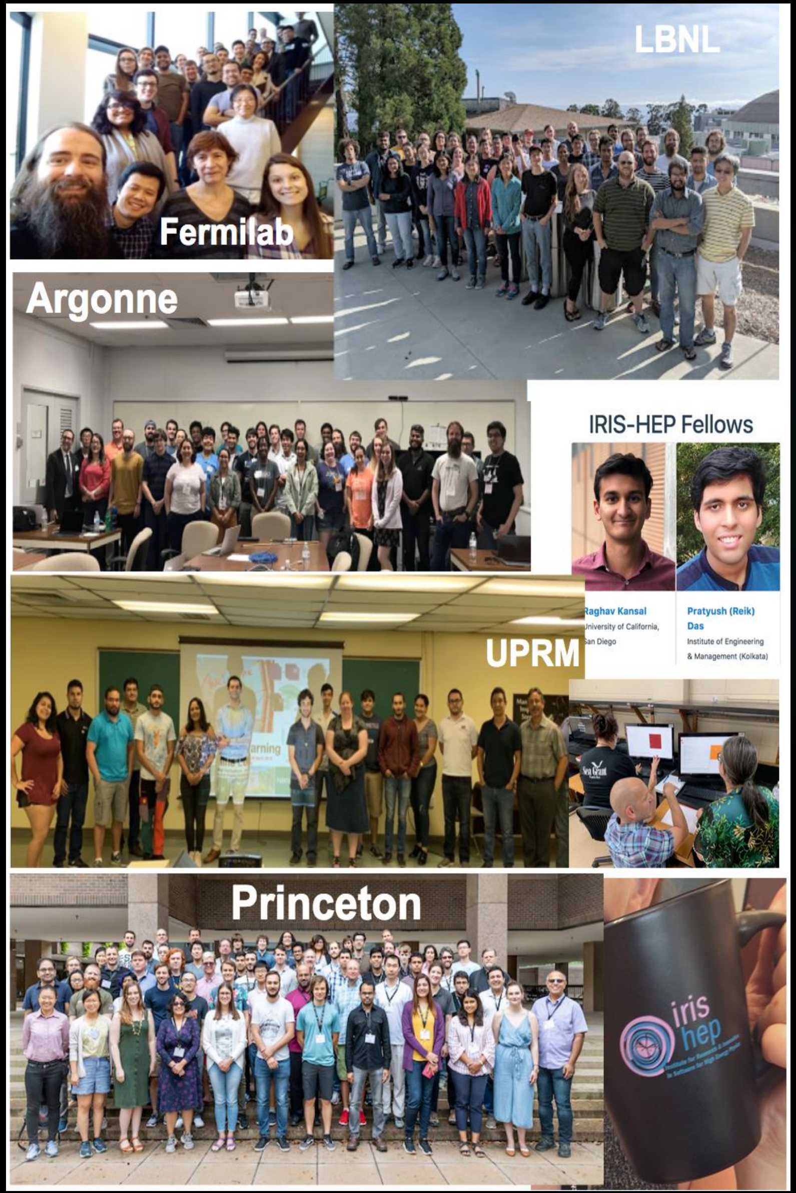

Management

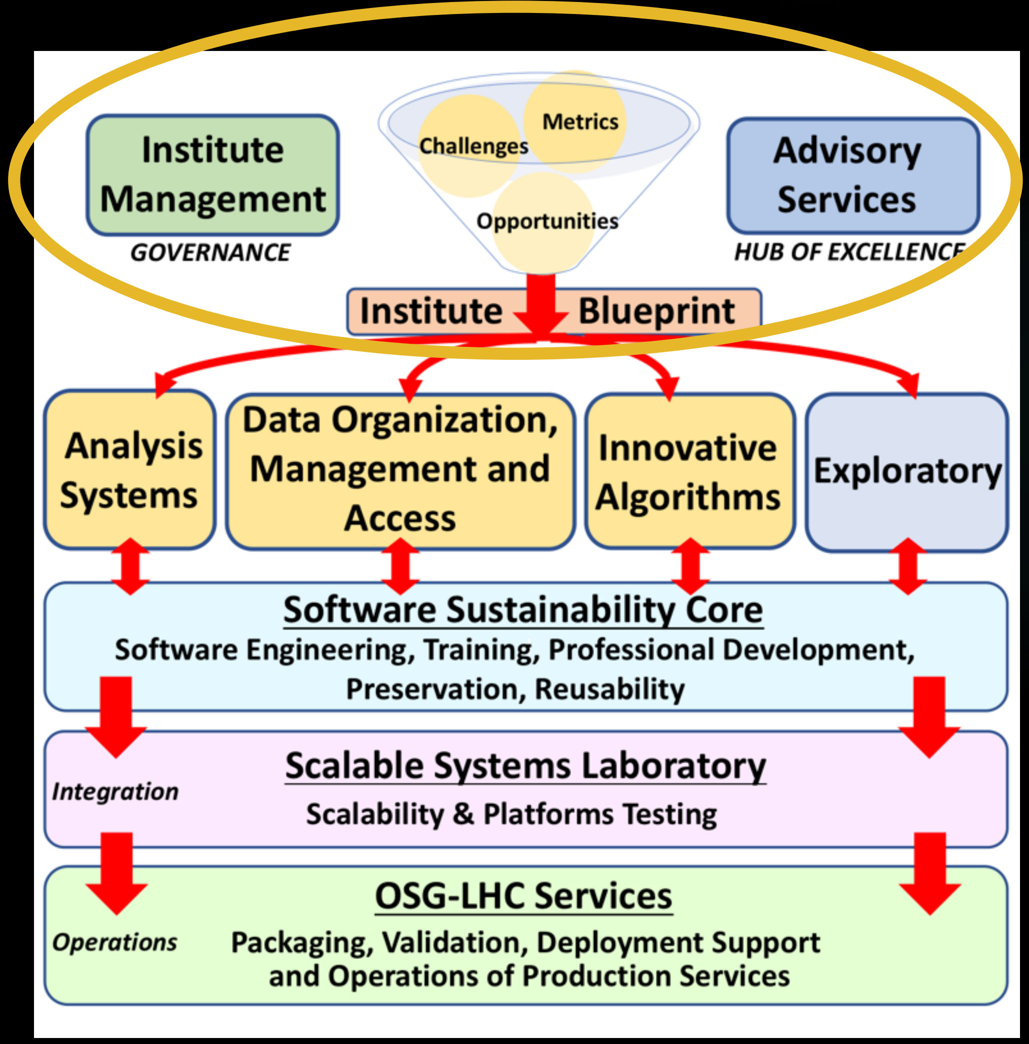
Institute Management
GOVERNANCE
Challenges
Metrics
Opportunities
Advisory Services
HUB OF EXCELLENCE
Institute Blueprint
Analysis Systems
Data Organization, Management and Access
Innovative Algorithms
Exploratory
Software Sustainability Core
Software Engineering, Training, Professional Development, Preservation, Reusability
Integration
Scalable Systems Laboratory
Scalability & Platforms Testing
Operations
OSG-LHC Services
Packaging, Validation, Deployment Support and Operations of Production Services

Executive Board

- Coordinates day-to-day operations
- Responsible for the near-and-long term schedule.
  - E.g. agency reviews & reports, annual meetings, Snowmass, etc.
- Coordinates direction and funding of the institute.
- Meets weekly.
- Include members of the US LHC experiments
  - Part of making sure work is grounded in what the experiments and field need

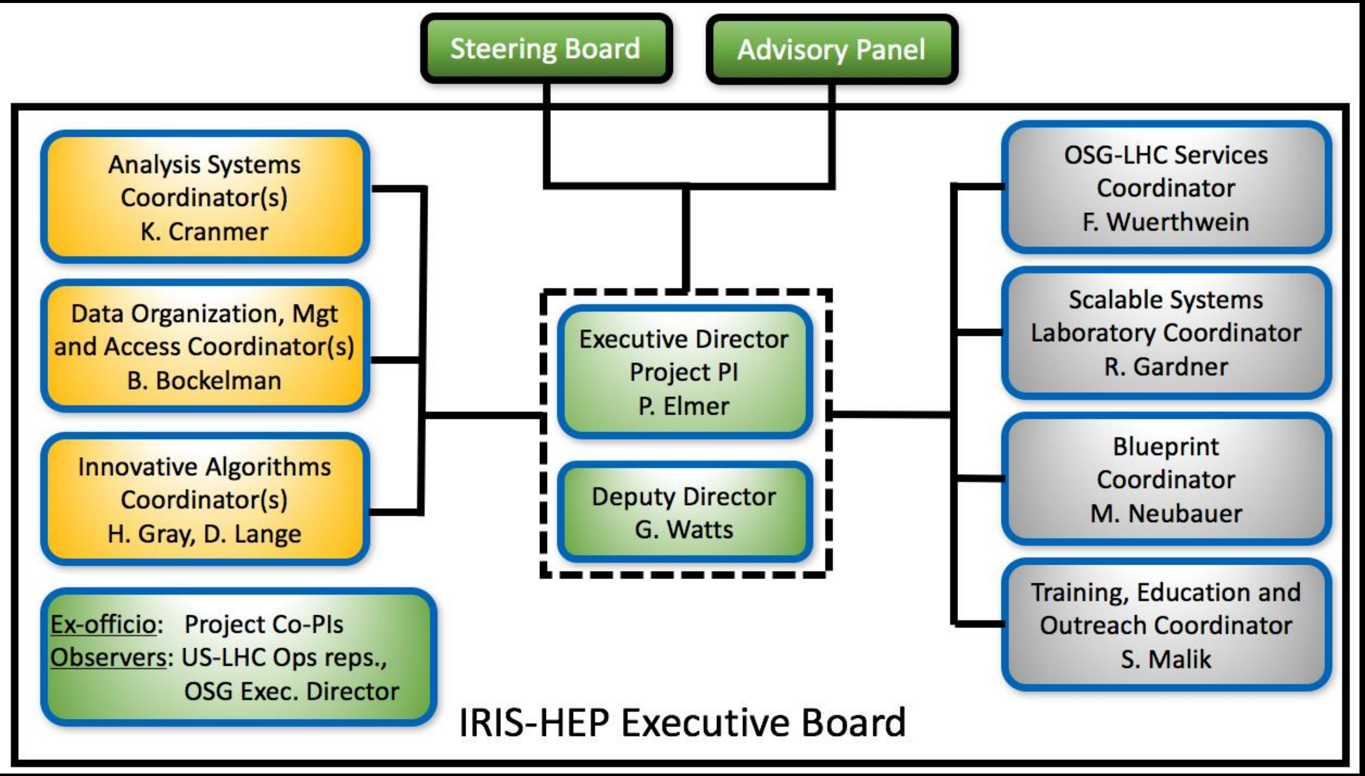

Steering Board & Advisory Panel

- Steering Board
  - Made up of stakeholders in the Institute's work
    - US ATLAS/CMS, ATLAS/CMS, LHCb, HSF, WLCG
  - Meetings quarterly
  - Feedback on direction
  - Place for experiments to discuss work and adoption (or not).
- Advisory Panel
  - Make up of non-stakeholder members
  - Global field perspective
  - Feedback on field-strategy.

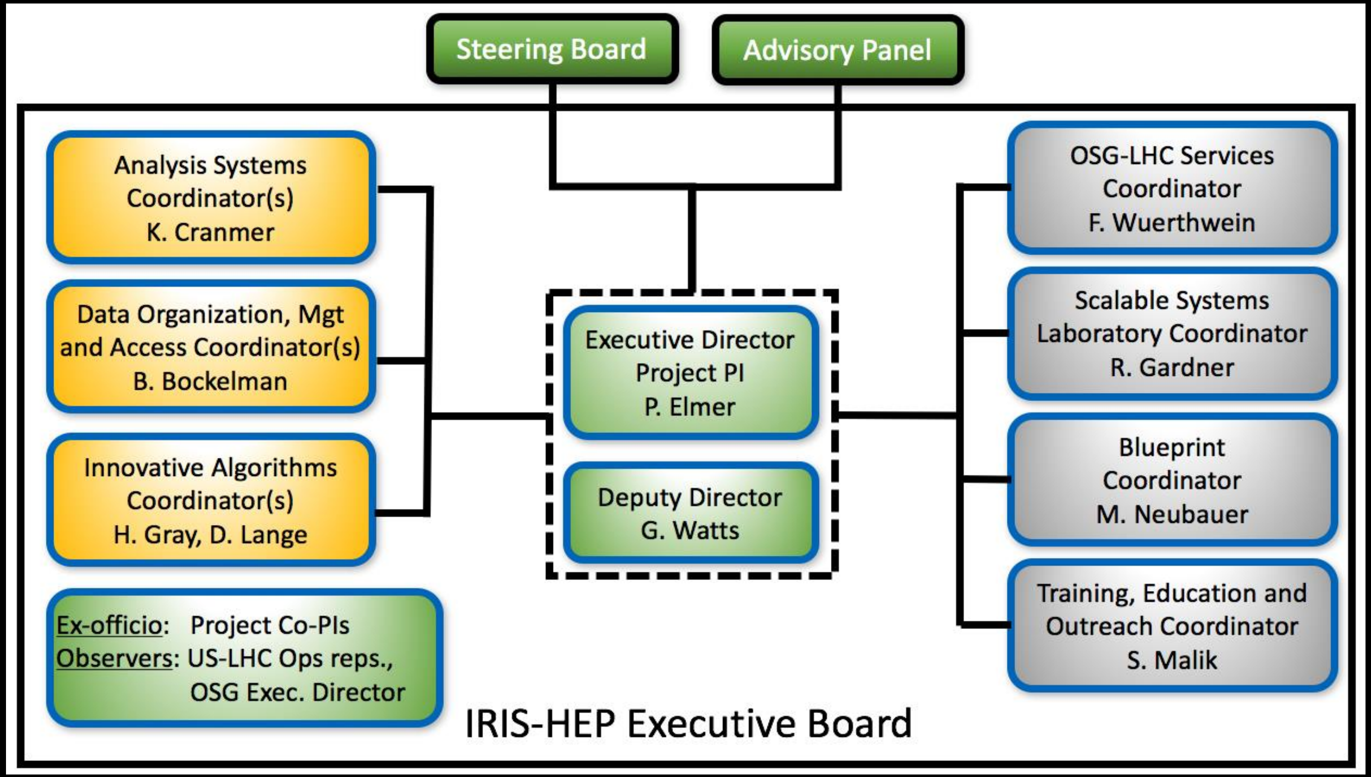

Milestones & Metrics

Label	Management Description	Type (M/D)	Y1Q1	Y1Q2	Y1Q3	Y1Q4
G1.1	Organize and execute the IRIS-HEP kickoff workshop	M				
G1.2	Establish initial web presence	D				
G1.3	Establish initial community and team mailing lists	D				
G1.4	Complete the design-phase project execution plan	D				
G1.5	Execute on Monthly and Quarterly reporting responsibilities	D				
G1.6	Finish execution of all Year 1 subawards	M				
G1.7	Advisory Panel and Steering Board membership finalized	M				
G1.8	Project support staff (project office, project manager) hired	M				
G1.9	Host Advisory Panel meetings	M				
G1.10	Host quarterly Steering Board meetings	D				
G1.11	Complete project staffing	M				
G1.12	Update and document Year 2 plans	M				
G1.13	First presentations at key conferences/workshops	M				
G1.14	First IRIS-HEP fellows	M				
G1.15	First Blueprint Workshops	M				
G1.16	Execute Year 2 subawards	M				
G1.17	Complete the execution-phase project execution plan	D				
G1.18	Organize and execute the IRIS-HEP general workshop	M				
G1.19	First publications from IRIS-HEP	M				

Blueprint Activity - Maintaining a Common Vision

- Small **"blueprint" workshops** 3-4 times per year with key personnel and experts
- Facilitate effective collaborations by building and maintaining a common vision
- Answer specific questions within the scope of the Institute's activities or within the wider scope of HEP software & computing.
- 21 Jun - 22 Jun, 2019 - Blueprint: Analysis Systems R&D on Scalable Platforms (*NYU*)
- 10 Sep - 11 Sep, 2019 - Blueprint: Accelerated Machine Learning and Inference

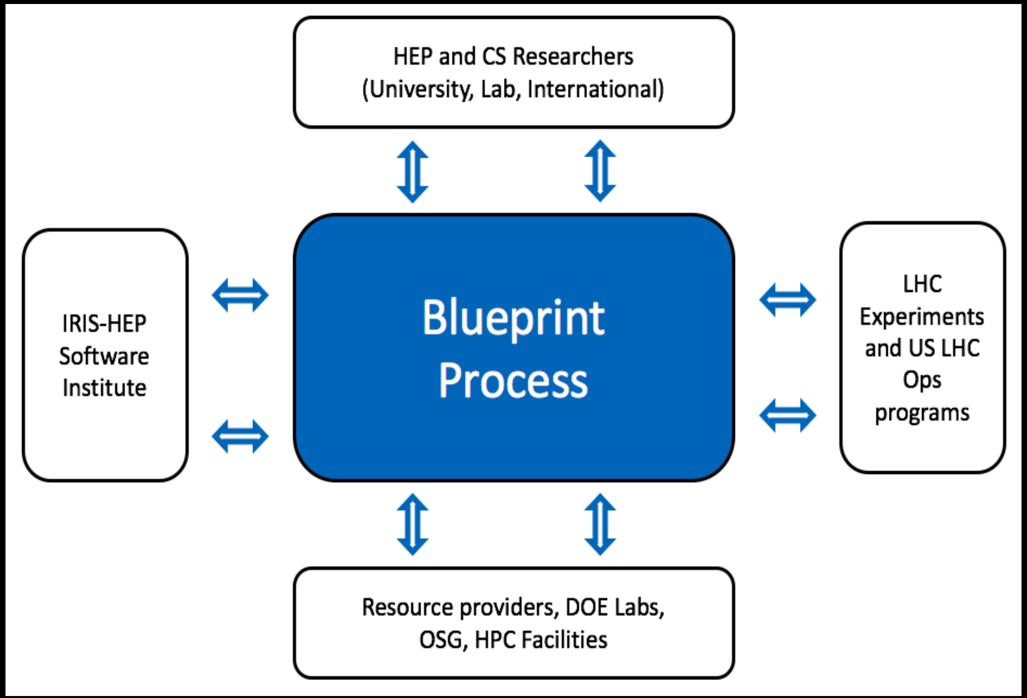

- 23 Oct - 25 Oct, 2019 - Blueprint: A Coordinated Ecosystem for HL-LHC Computing R&D (*Catholic University of America, Washington DC*)
- Others (e.g. Training) in planning

Highlight: A Coordinated Ecosystem for HL-LHC Computing R&D

How does the ensemble of US Software R&D efforts fit together to implement the HL-LHC Software/Computing roadmap described in the Community White Paper and meet the challenges of the HL-LHC? Which areas are not covered by US R&D efforts?

How do the US Software R&D efforts collaborate with each other and with international efforts? How do these efforts align with and leverage national exascale, national NSF OAC priorities and trends in the broader community?

How should the US R&D efforts be structured and organize in order to impact planned updates (all in ~2021/2022) to the HSF Community White Paper, the software/computing part of the US Snowmass process and HL-LHC experiment-specific software/computing TDRs?

~40 participants from US National Labs, the two US-LHC Operations Programs, the IRIS-HEP team, and NSF & DOE program directors/managers. [The meeting was organized with the involvement and blessing of the DOE program managers.]

This built on a similar workshop at CUA which we organized in 2017 during the conceptualization phase for IRIS-HEP.

Checkpoint the relationship of the IRIS-HEP efforts and collaborations with the two US-LHC Operations programs.

Explore how IRIS-HEP interacts with other NSF efforts and planned and ongoing DOE efforts (SciDAC, LDRD, ECP, ESNET, HEP-CCE, facilities, etc.)

Discussion-oriented: a report will be issued before Xmas.

Foster an Inclusive, Collaborative Environment

https://iris-hep.org/about/code-of-conduct

iris hep

About Connect Activities Jobs

IRIS-HEP Code Of Conduct

IRIS-HEP is a software institute funded by the National Science Foundation. It serves as a center for software R&D, functions as an intellectual hub for the larger community-wide software R&D efforts, and transforms the operational services required to ensure the success of the HL-LHC scientific program. Its success depends on its ability to engage a community from several institutions and disciplines with diverse skills, personalities and experiences.

Naturally, this implies diversity of ideas and perspectives on often complex problems. Disagreement and healthy discussion of conflicting viewpoints is welcome: the best solutions to hard problems rarely come from a single angle. But disagreement is not an excuse for aggression: humans tend to take disagreement personally and easily drift into behavior that ultimately degrades a community. This is particularly acute with online communication across language and cultural gaps, where many cues of human behavior are unavailable. We are outlining here a set of principles and processes to support a healthy community in the face of these challenges. Fundamentally, we are committed to fostering an open, productive, harassment-free environment for everyone. Rather than considering this code an exhaustive list of things that you can't do, take it in the spirit it is intended - a guide to make it easier to enrich all of us and the communities in which we participate. Importantly: as a member of our community, you are also a steward of these values. Not all problems need to be resolved via formal processes, and often a quick, friendly but clear word on an online forum or in person can help resolve a misunderstanding and de-escalate things. However, sometimes these informal processes may be inadequate: they fail to work, there is urgency or risk to someone, nobody is intervening publicly and you don't feel comfortable speaking in public, etc. For these or other reasons, structured follow-up may be necessary and here we provide the means for that: we welcome reports by contacting any member of the IRIS-HEP Executive Board. This code of conduct applies equally to all community members in all Institute situations online and offline, including conferences, training events, mailing lists, forums, GitHub organizations, chat rooms, social media, social events associated with conferences and events, and one-to-one interactions.

By embracing the following principles, guidelines and actions to follow or avoid, you will help us make IRIS-HEP a welcoming and productive community. Feel free to contact any member of the IRIS-HEP Executive Board.

Not a collection of isolated Projects

Towards A Common Goal

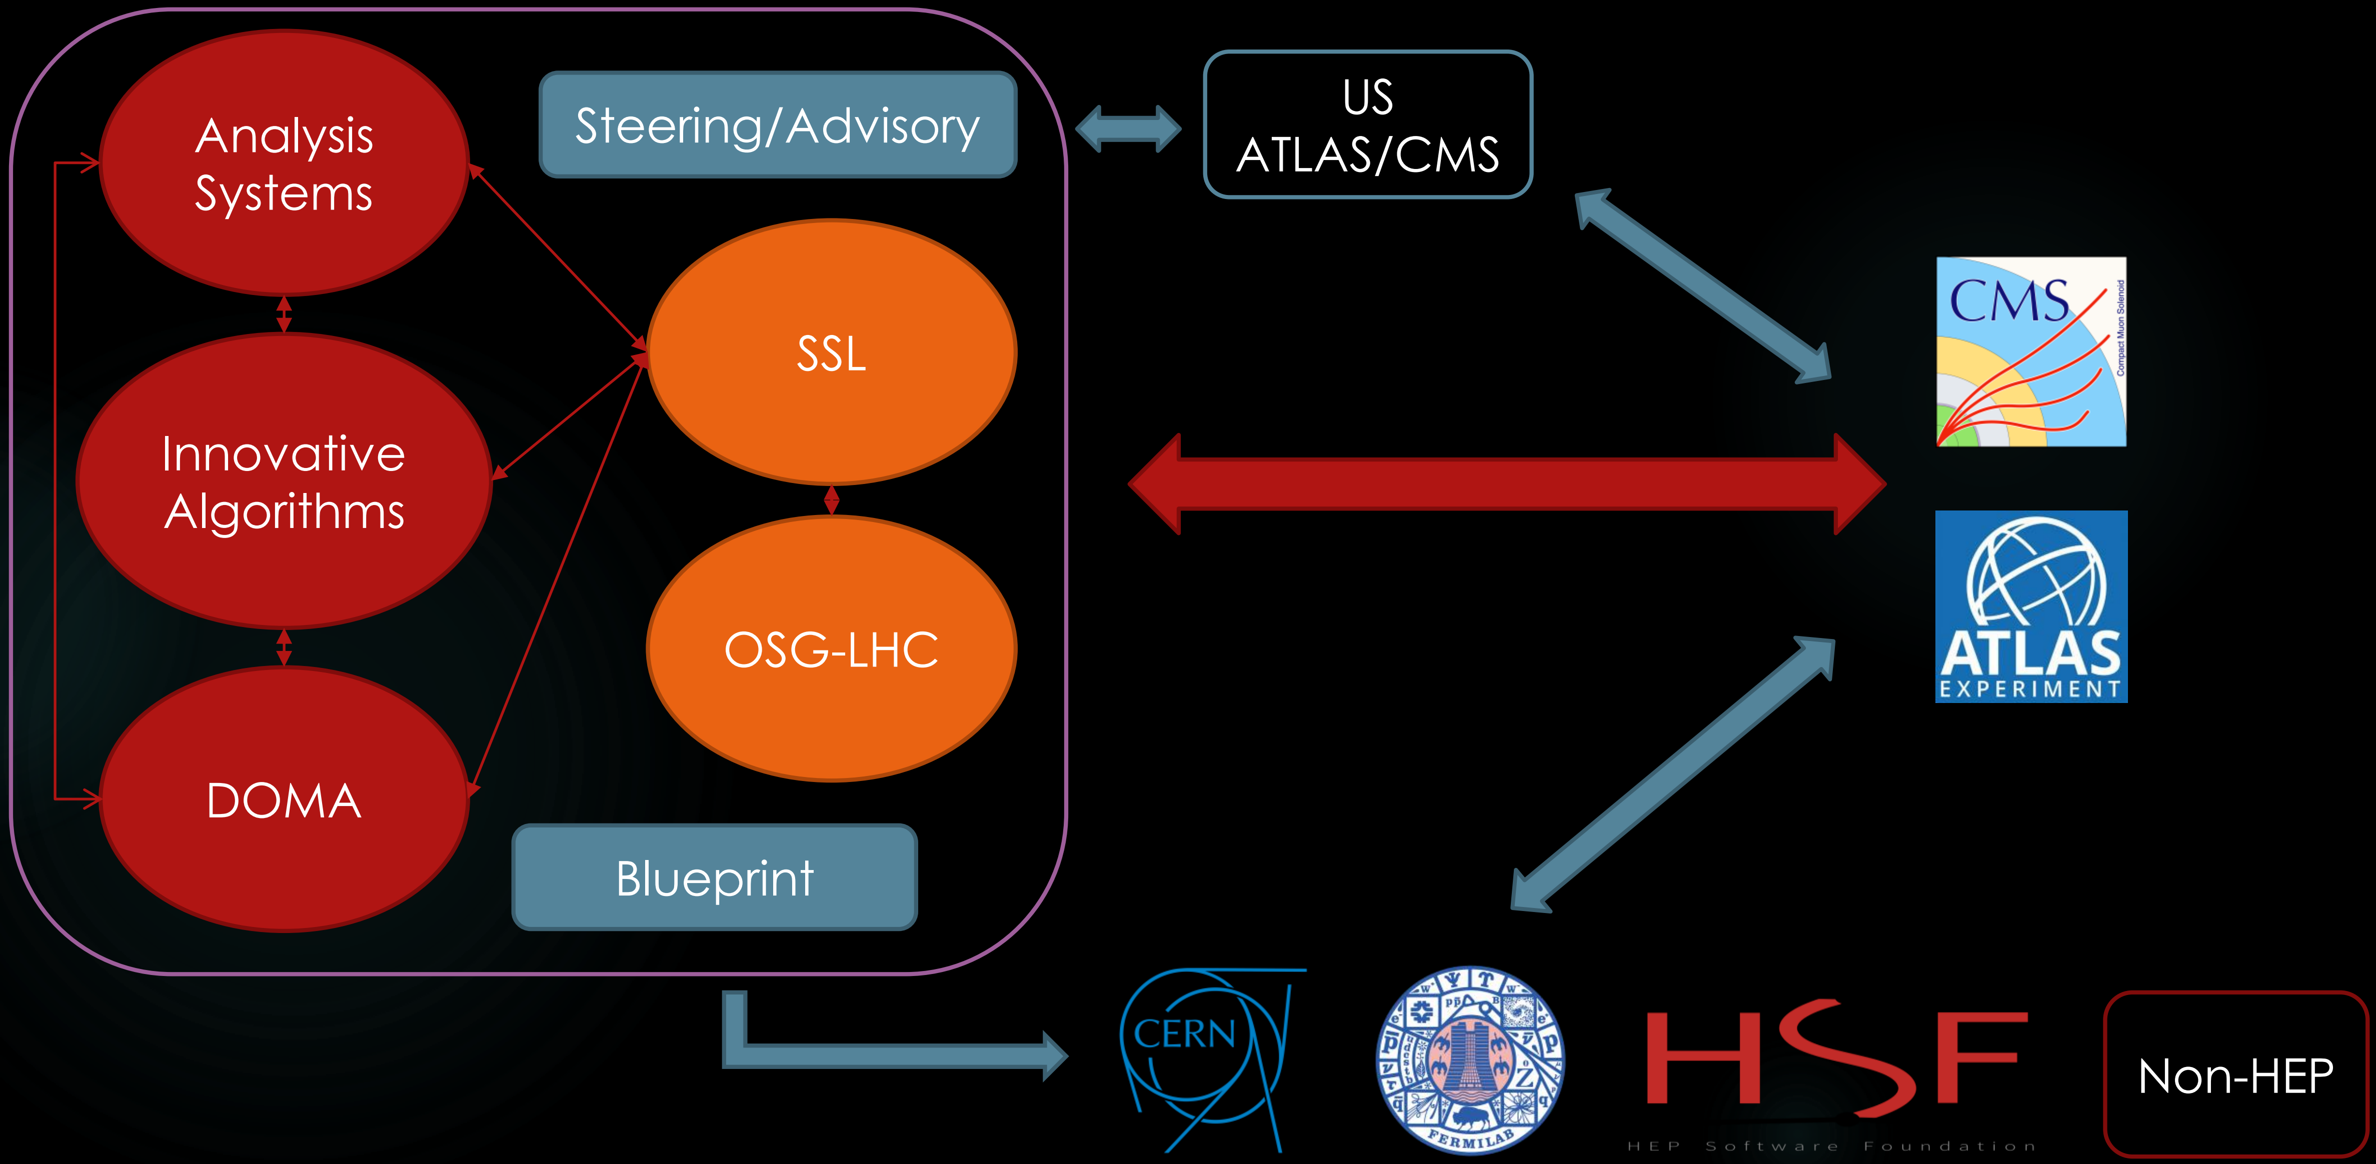

Conclusions

- A tightly knit collection of university efforts
- Interactions are critical to the success of the project
  - Keeping the university efforts aligned towards a broader common goal (like Analysis Systems)
  - Keeping the efforts of the project aligned with the physics and computing goals of the LHC experiments
  - Interacting with outside community to foster idea and software exchange
- Play to the US University Group's Strengths
  - The groups that went into IRIS-HEP helped choose the topics from the CWP
  - All done in the context of the CWP
- Vision of what computing looks like in the HL-LHC era
  - This will include regular updates of the CWP
  - Use this vision to drive evolution of IRIS-HEP as we let go of things we try and push things the work.