IRIS-HEP and Opportunities for Collaboration

G. Watts (for IRIS-HEP)
March 18, 2024

IRIS-HEP

"IRIS-HEP is a software institute funded by the National Science Foundation. It aims to develop the state-of-the-art software cyberinfrastructure required for the challenges of data intensive scientific research at the High Luminosity Large Hadron Collider (HL-LHC) at CERN, and other planned HEP experiments of the 2020's."

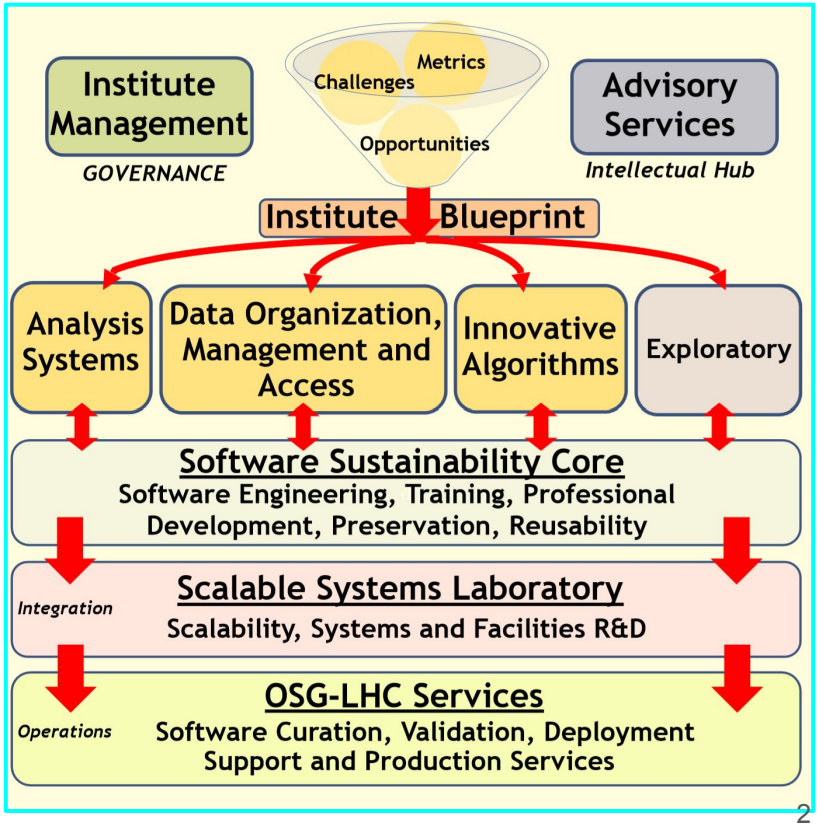

IRIS-HEP

Gaps:

1. Raw resource gaps
2. Scalability of the distributed computing cyberinfrastructure,
3. Executing analyses at the HL-LHC scale
4. Sustainability of the software and computing system through the lifetime of the HL-LHC and beyond

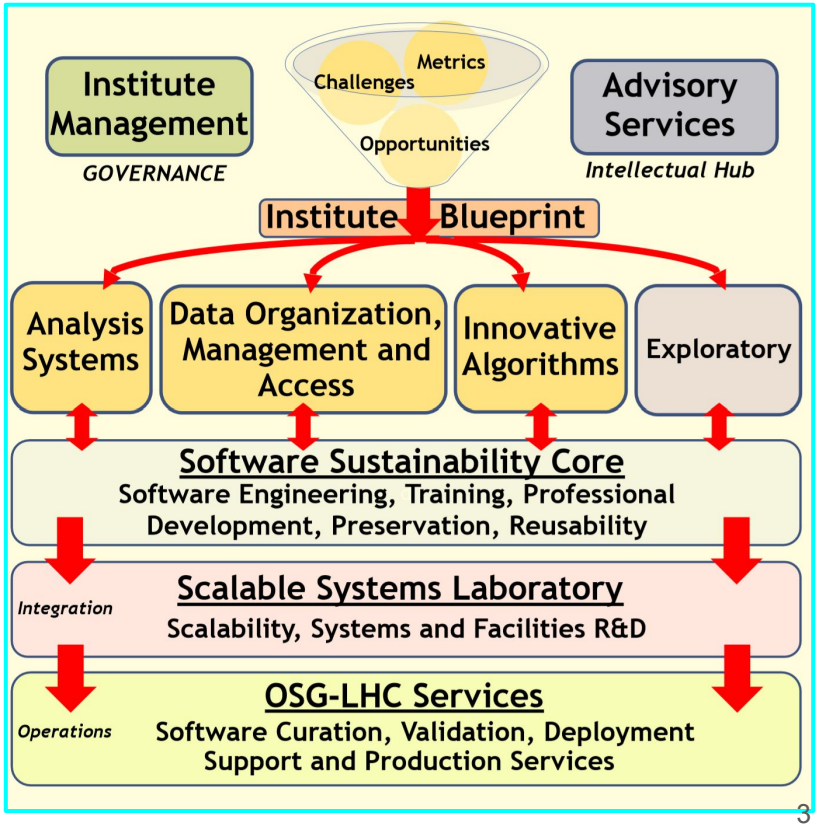

From the Goals/Current Focus Talk Earlier…

NPPS SDCC

ATLAS Distributed Computing: Towards HL-LHC

The flagship: PanDA workload manager

- Developed by BNL and UT Arlington (UTA) 2005-present
- 24x365 processing on ~800k cores globally for ATLAS
- All workloads, production and analysis, at ~150 facilities, ~1500 users, ~1M jobs/day, Exabyte throughput/year
- All resource types: WLCG, clouds (GCP, AKS, k8s-based clusters), HPCs, BOINC, …

Production and Distributed Analysis (PanDA) workload management system

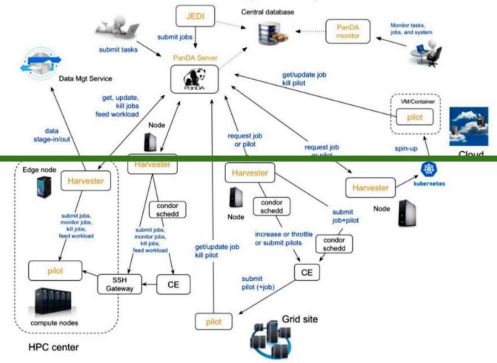

Built to scale smoothly to HL-LHC (2030+)

- Efficient optimization of workflows to economize storage and processing at the HL-LHC
- Basis, with Rucio, of the Data Carousel

Emphasis now on analysis functionality

- **Ease of use:** Improved interactivity through Jupyter and **interactive** PanDA monitor extensions
- **Speed:** Latency reduction throughput the system, more direct message-driven internal communication
- **Effective resources:** PanDA queues and sites leveraging non-grid platforms like **k8s**, **clouds**
- **Complex workflows:** PanDA as an **engine for large scale AI/ML** and other complex workflows

ATLAS and CMS CPU needs for HL-LHC

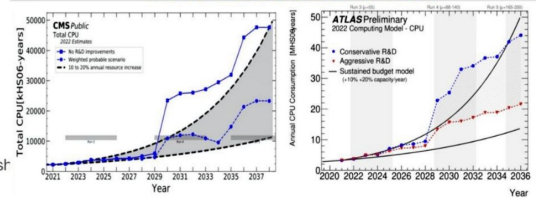

Brookhaven National Laboratory

A. Klimentov, T. Wenaus

Worksh

Very similar language when we talk about high level goals for IRIS-HEP:

- Time to insight
- Ease of use
- Solving the “gaps” between us an HL-LHC computing

IRIS-HEP

Focus Areas:

- Analysis Systems - HL-LHC scale analyzes, high-data-rate columnar analysis, UX, integration of ML into the pipeline
- DOMA - exabyte-scale production dataset management and delivery to analysis facilities, scaling bulk data, modernization, and data delivery.
- Innovative Algorithms - real-time data processing trigger and offline, focus on tracking.

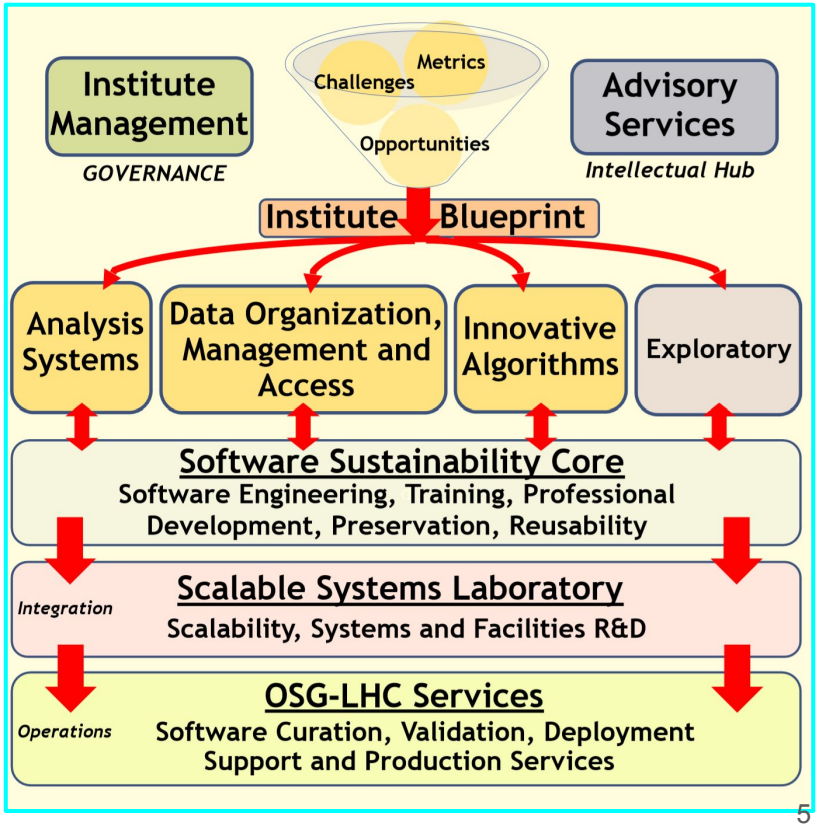

IRIS-HEP

Services

- Scalable Systems Laboratory (SSL) - k8s devops platform for facilities R&D, support for training, summer schools, and development (coffee-casa/ServiceX)
- OSG - Services for the LHC, production services for LHC's cyberinfrastructure

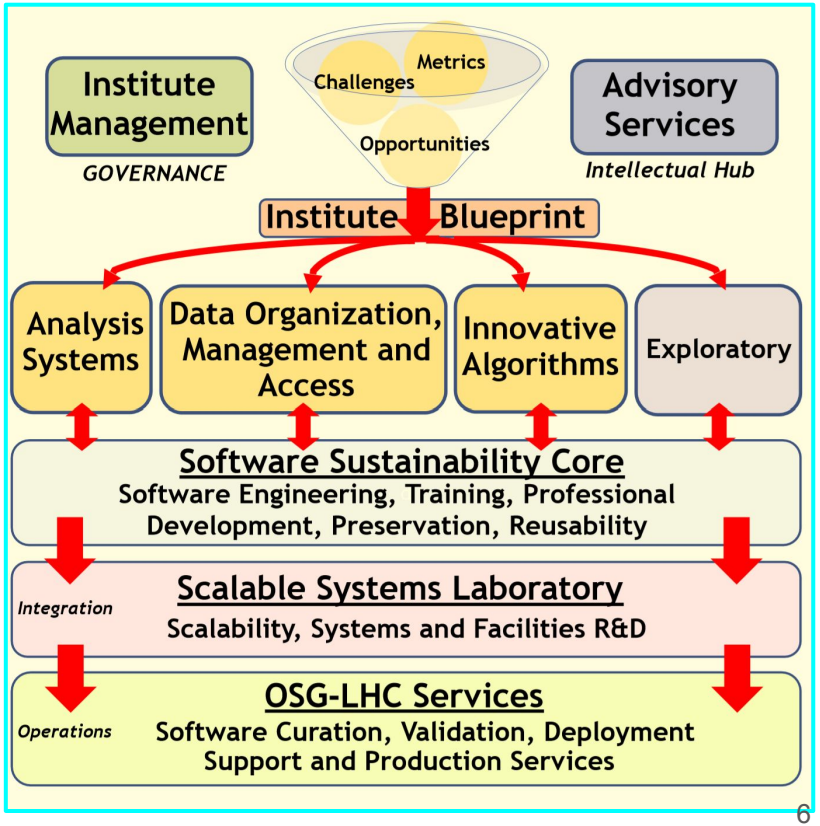

IRIS-HEP

Intellectual Hub

- Work with other large projects around the world
- Host a large slack for everyone interested in collaborating (with us, or with each other on these matters)
- Host Blueprint Meetings
  - Building consensus in the community (outward)
  - Building a strategic plan (inward)

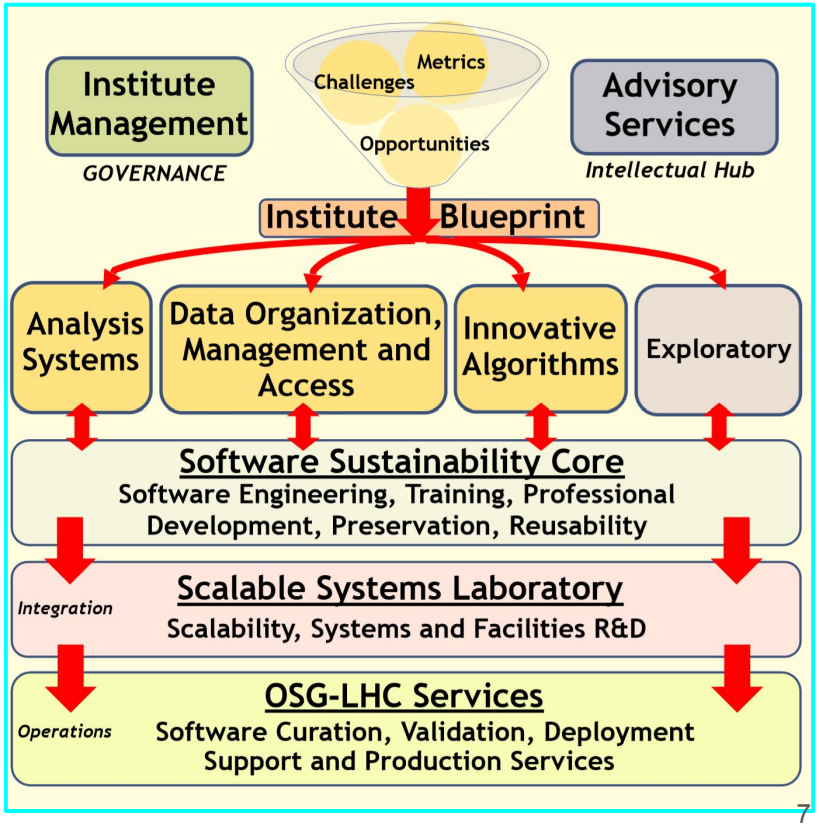

IRIS-HEP

Training

- Centered around a sustainable workforce
- Scaling the training to reach a large fraction of US graduate students
- Work hand-in-hand with the HSF
- Aim for field **sustainability** through depth in software knowledge.

Full details during Lightning Talk by Villanueva later today.

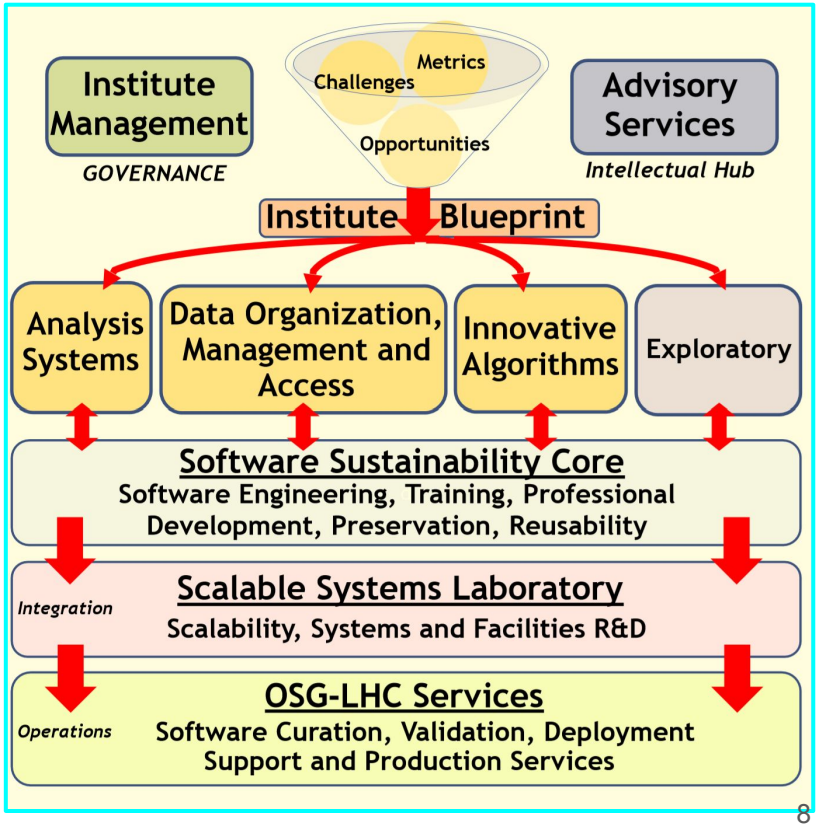

IRIS-HEP

Community Contact(s)

- Informal: members are part of the community, embedded in experiments, etc.
- Projects mostly hosted on github
- Steering Board: Formal board of stakeholders (experiments, US Operations programs, etc.)
- Executive Board: Leaders of all the major parts of IRIS-HEP
- Peter Elmer, Brian Bockelman, & Myself for direct queries at

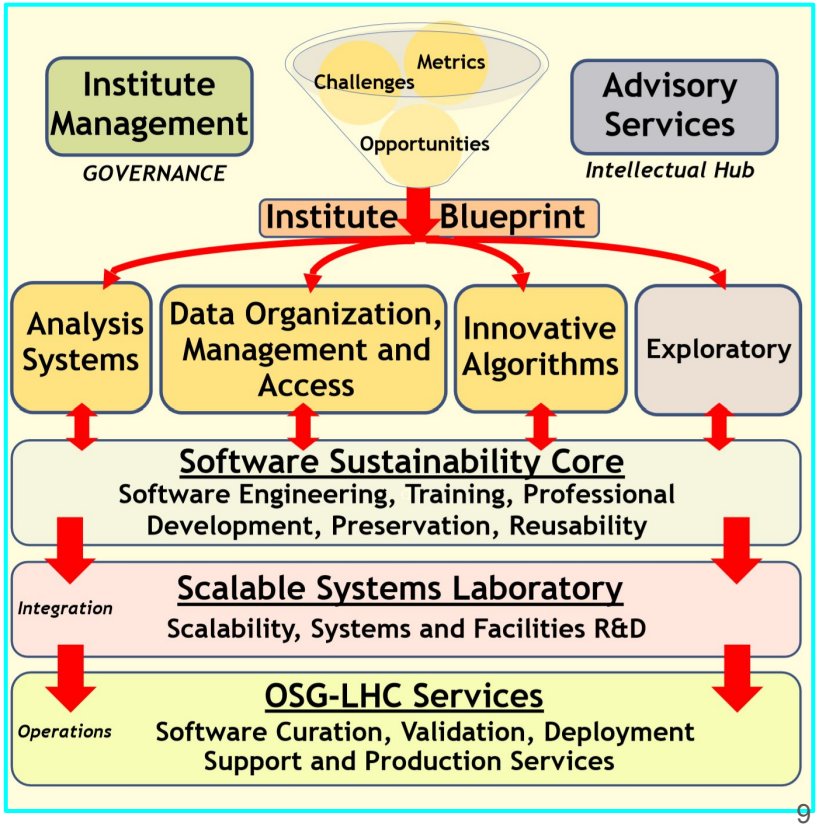

Projects

A list of **all** projects!

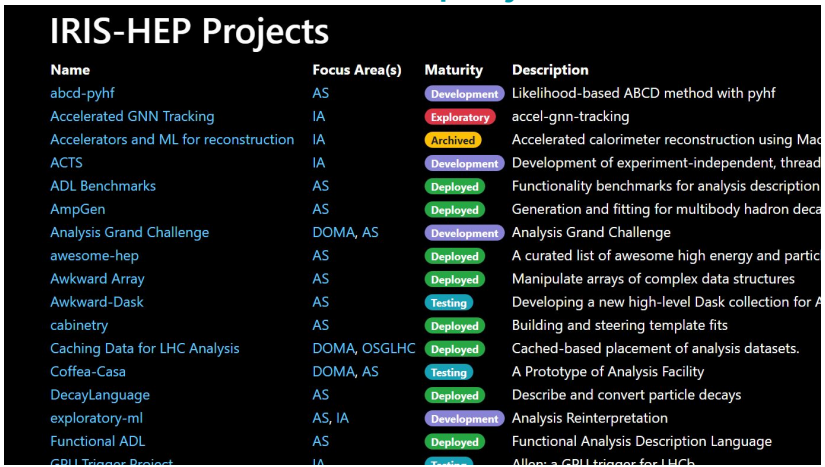

Our website is an index of almost everything we do

- We use it ourselves to track progress
- We draw our reports from updates to the web site
- It is quite up to date as a result!

Also listings per-area

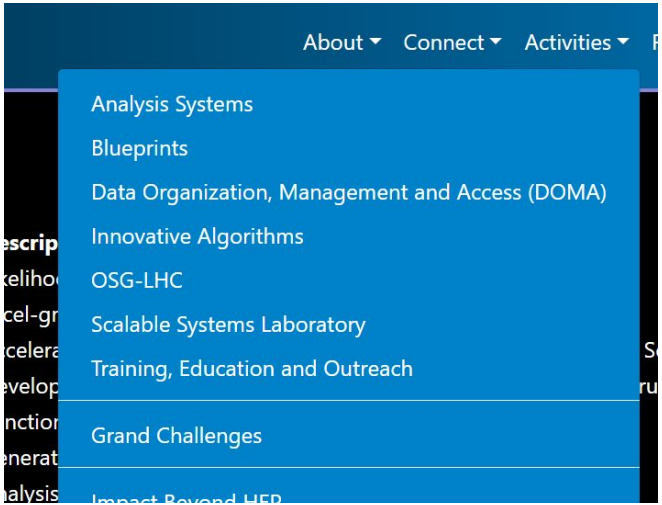

Analysis Grand Challenge

“Time to insight”
“Efficient use of resources”

The “AGC”

Goal: Show the analysis pipeline is integrated and scales efficiently.

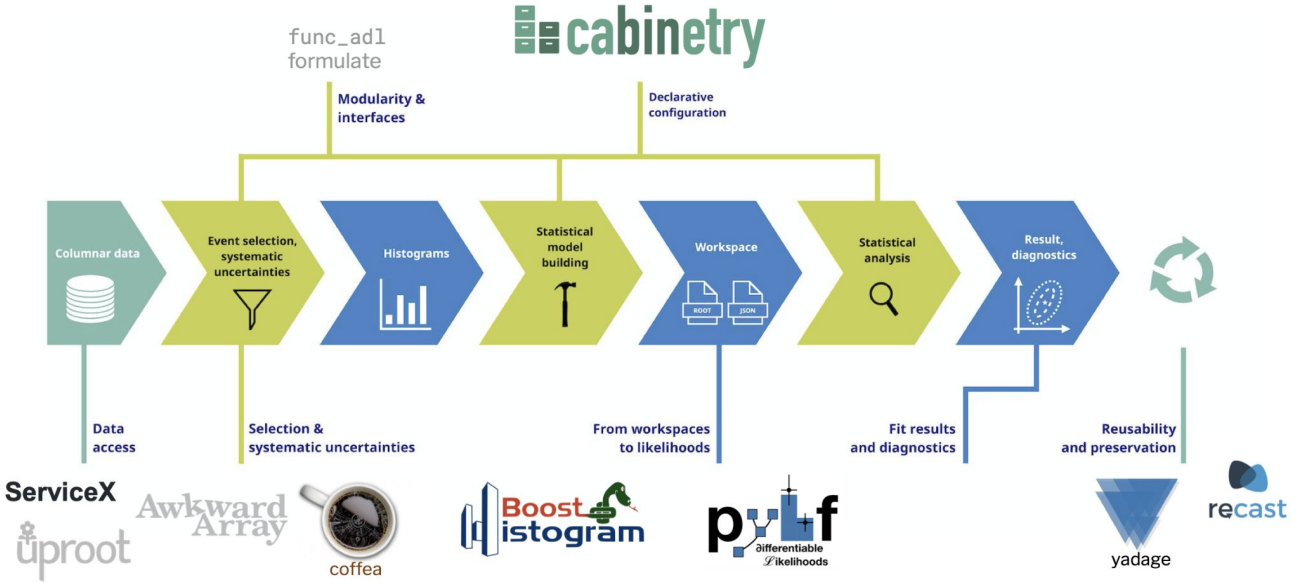

Requires:

- Physicist
- Software
- Facility

To all work in concert

The integrated nature of this is new

Scale, scale, scale…

~Presentations at: SciPy 2024

IRIS-HEP and parts of the python stack…

- ServiceX - data delivery
- Uproot - loading TTree, RNTuple, THist data
- Awkward array & friends - numpy for jagged data
  - Vector, etc. - 4 vector operations
  - PHYSLITE schema (w/ATLAS) - making it possible to use!
  - Dask-awkward - operation-level scaling (not MPI)
- Cabinetry - Fitting workspace (think TRexFitter)
- Pyhf - HistFitter, using GPU's
- hist/boost-histogram - modern multithreaded histograms
- Scikit-hep - Community project of packages (e.g. start here)

Benchmarks To Software and Facilities

tt CMS Open Data Benchmark Analyses (Prelim by IRIS-HEP)

Internal Benchmark Analyses (Defined by ATLAS/CMS, LHCC) (future)

Software Repo (docs)

Bottlenecks in Software and Hardware

- Distributed (gross via coffea processors)
- Distributed (dask)*
- ML (local)
- ML (triton & similar)
- Systematic Errors
- ServiceX

* near future

L. Gray’s talk at ACAT/SB last Week

Future future

- Differentiable programmings
- GPU’s (etc.) - pyhf now

Private version for ATLAS internal use? Sure!

Analysis Facilities

What is an Analysis Facility?

White Paper (from HSF)

- Good survey of capabilities currents ones have
  - And what new ones might be beneficial
- Community is feeling its ways

Question:

- Document lays out many possible paths forward
- E.g. Silo’d services or integrated services? Both?

Analysis Facilities White Paper

Diego Ciangottini [1,b], Alessandra Forti [2,b], Lukas Heinrich [3,b], Nicola Skidmore [4,b], Eduardo Rodrigues [5], Graeme A. Stewart [6], Gordon Watts [7], Mark S. Neubauer [8], Doug Benjamin [9], Oksana Shadura [10], Andrès Pacheco Pages [11], Antonio Delgado Peris [12], Nick Smith [13], Jonas Eschle [14], Liz Sexton-Kennedy [13], Evangelos Kourlitis [3], Alexander Held [15], José Flix Molina [12], Antonio Perez-Calero Yzquierdo [12], Burt Holzman [13], Niladri Sahoo [16], Kevin Pedro [13], Daniele Spiga [1], Jamie Gooding [17], Giordon Stark [18], Clemens Lange [19], Piergiulio Lenzi [1], Thomas Kuhr [3], Cristiano Alpigiani [7], Verena Martinez Outschoorn [20], Dmitry Kondratyev [21], Stefan Piperov [21], Brij Kishor Jashal [22, 23], Robert Gardner [24], Ilija Vukotic [24], Fengping Hu [24], Lincoln Bryant [24], Lindsey Gray [13], José Hernández [12], Brian Bockelman [25],

b editor, 1 INFN, 2 University of Manchester, 3 Technische Universität München, 4 University of Warwick, 5 University of Liverpool, 6 CERN, 7 University of Washington, 8 University of Illinois at Urbana-Champaign, 9 Brookhaven National Laboratory, 10 University Nebraska-Lincoln, 11 IFAE, 12 CIEMAT, 13 FNAL, 14 Syracuse University. 15 University of Wisconsin–Madison, 16 University of Birmingham, 17 Technische Universität Dortmund, 18 SCIPP UC Santa Cruz, 19 Paul Scherrer Institute, 20 University of Massachusetts Amherst, 21 Purdue University, 22 IFIC, 23 TIFR, 24 University of Chicago, 25 Morgridge Institute for Research,

Abstract

This whitepaper presents the current status of the R&D for Analysis Facilities (AFs) and attempts to summarize the views on the future direction of these facilities. These views have been collected through the High Energy Physics (HEP) Software Foundation’s (HSF) Analysis Facilities forum, established in March 2022, the Analysis Ecosystems II workshop, that took place in May 2022, and the WLCG/HSF pre-CHEP workshop, that happened in May 2023. The paper attempts to cover all the aspects of an analysis facility.

IRIS-HEP’s Coffea-Casa

Integrated solution

- Modular design that can share dependencies with other projects
- Log in once, access to everything (auth)
- JupyterHub is the main access
- k8s deployment, configurable for different sized facilities
- Deployed at UChicago (ATLAS) and Nebraska (CMS)
  - ATLAS IAM sign in is fully functional, but there is still on-going work to enable access to ATLAS specific datasets

coffea-casa

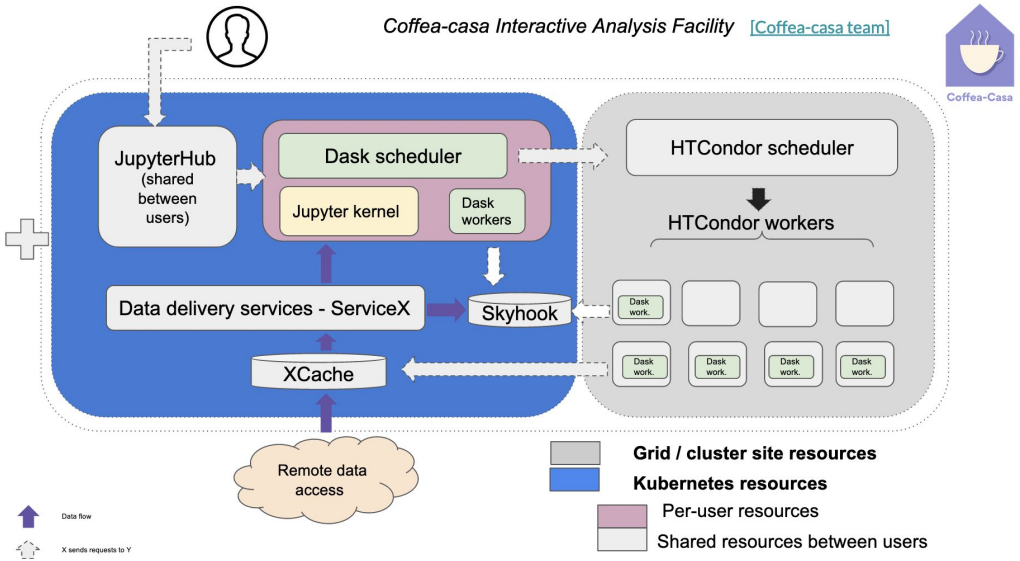

ServiceX

Columnar Data Delivery System

- ROOT files or parquet files
- Over the web, cached
- Inside AF no need to download
- k8s based, moving to OKD

Current Backend Work

- Preparing for next push
- Cross-facility integration
- Scaling and sharing resources on cluster with other services

ServiceX

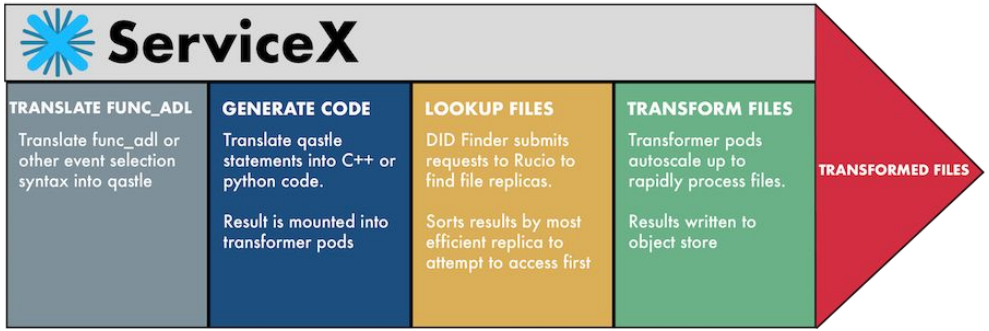

- 10 TB
- ~1000 Cores
- 30% of file read
- 100 columns
- 30 minutes

Scalable Systems Lab (SSL)

UChicago SSL Effort

- Platform for IRIS-HEP service development (DOMA, Analysis Facilities) and integration with LHC production environments
- R&D on evolving computing facilities based on flexible foundation with cloud-native tools (K8s)
- Effort aligned with US ATLAS Facilities R&D program
  - Example: Joint ATLAS / IRIS-HEP Kubernetes Hackathon (24-26 April 2024)

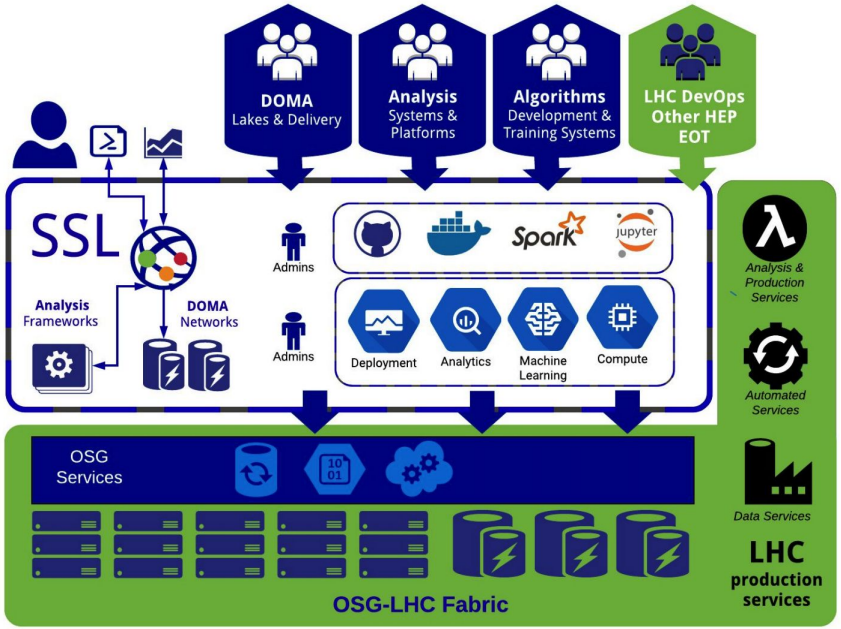

Training

Early work on pulling **new people** into the field/pipeline.

Non-HSF Efforts

Fellows Program

- Similar to Google Summer of Code
- We have funds for undergraduates over the summer
- Have a project aligned with IRIS-HEP goals? Need an undergraduate? We might be able to help
- Currently have open a large number of projects.

See talk later today for HSF related work

HSF
HEP Software Foundation

Full details during Lightning Talk by Villanueva later today.

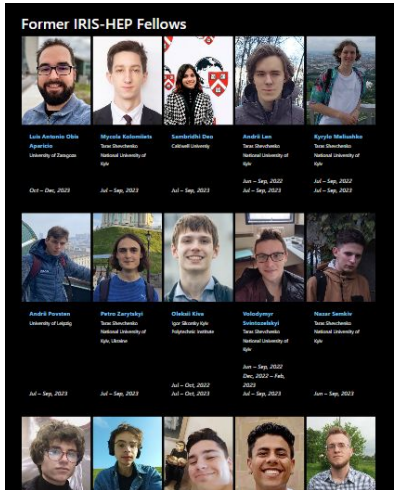

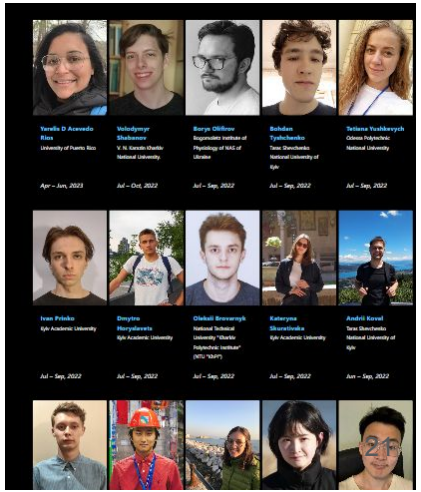

Training Events

US ATLAS and US CMS Summer Meetings

- 2-day training programs associated with the summer meetings (US ATLAS, US CMS)
- Occur after the meetings
- Supply funds to get students on-location and board for the 2 nights
- Advanced topic range (programming on GPU's, columnar analysis, statistical models)
- In addition to experiment specific training that occurs during summer meeting

SSL Binder (River, Flatiron, National Research Platform) https://binderhub.ssl-hep.org and ***K8s stretched cluster resources at US ATLAS Tier2s and at UVictoria***

OSG

Overview of OSG-LHC

OSG-LHC provides:

- An **integrated software stack**, coordinated with the global community.
  - *Software lifecycle management, including ingestion, releases, “orphanage”, and retirement process.*
- A **cybersecurity context**
  - *From policies to operational security*
- Operations
  - ***Accounting and Runtime environment***
- **Network performance monitoring**
- Various **coordination functions** with the Worldwide LHC Computing Grid (WLCG), the US LHC Operations program, and other science communities.
- OSG-LHC staff works closely with WBS 2.3:
  - Reports regularly on OSG and HTCondor software updates at bi-weekly US ATLAS Computing Facilities meetings
  - Security operations: participates in CERN SOC hackathons, cybersecurity threat intelligence sharing
  - Network data collection and analytics with the community (weeky R&D meetings chaired by Shawn)
  - Currently working o:n K8s WLCG accounting for UTA, NET2 and Google projects; Token enabled Xrootd for SWT2 (UTA and OU)

Production enabled by the OSG

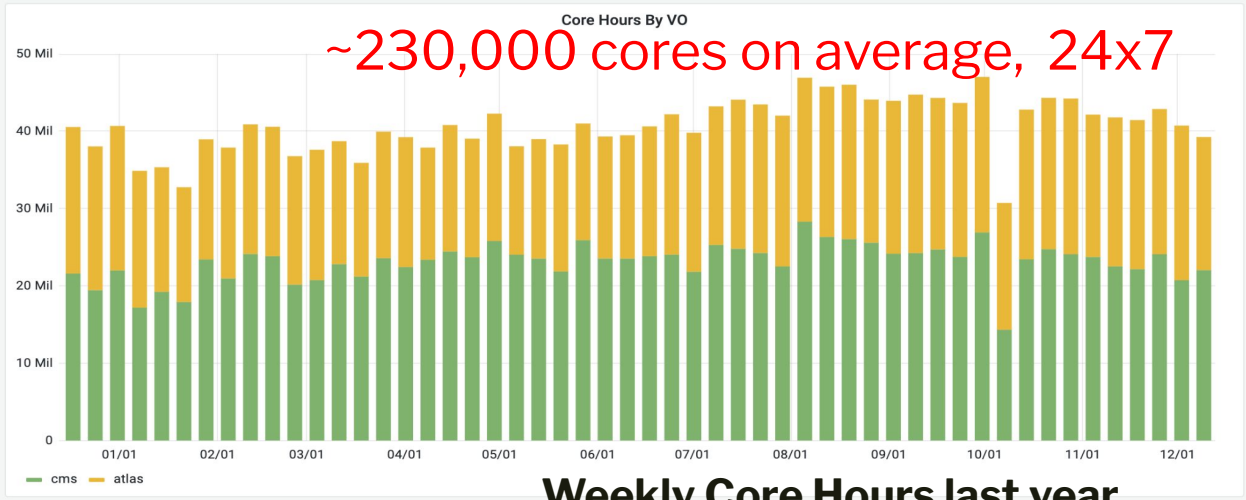

Weekly Core Hours last year

US LHC Operations Programs are the primary hardware providers, and thus customer of OSG. Lots of coordination activities, at many levels.

- Annual workshop of US LHC facilities programs co-located with OSG AHM
- OSG-LHC staff attending relevant ATLAS and CMS meetings in US & CERN
- US LHC Ops program management part of OSG management
- …. and many more ….

An integration example with OSG - ePIC/EIC

The Electron-Proton/Ion Collider (ePIC) Collaboration was formed to design, build, and operate the first experiment at the Electron-Ion Collider (EIC). Currently there are two main computational efforts and cyberinfrastructure integrations to BNL.

First:

- An EIC submitter at BNL to the JLab gWMS pool and OSPool
- An ePIC/EIC submitter to JLab and OSPool gWMS pools from JLab.
- Both submitters - Access Points - leverage the JLab FE to request pilots from the OSG gWMS factory and use two gWMS collectors deployed by JLab (one in HA mode). ePIC/EIC jobs run on both BNL and JLab (gracc)

Second:

- An investigation to deploy Harvester/PaNDA at BNL for EIC and send Harvester jobs to the JLab farm and the EIC pool.

Conclusions

Presentations and Publications

Presentations

Presentations: by area • by institution • by date • by person. Publications: by date Related: Open Science.

Presentations by the IRIS-HEP team

July, 2024

- 1 Jul 2024 - "A future presentation", Ana Peixoto, Some Workshop

March, 2024

- 11 Mar 2024 - "ServiceX, the novel data delivery system, for physics analysis", KyungEon Choi, ACAT 2024
- 11 Mar 2024 - "Awkward Family (poster)", Jim Pivarski, ACAT 2024
- 11 Mar 2024 - "Bridging Worlds: Achieving Language Interoperability between Julia and Python in Scientific Computing", Ianna Osborne, 22nd International Workshop on Advanced Computing and Analysis Techniques in Physics Research - ACAT2024
- 11 Mar 2024 - "Retrieval Augmented Generation for Particle Physics: A Case Study with the Snowmass White Papers and Reports", Gordon Watts, ACAT 2024
- 11 Mar 2024 - "From Amsterdam to ACAT 2024: The Evolution and Convergence of Declarative Analysis Language Tools and Imperative Analysis Tools", Gordon Watts, ACAT 2024
- 11 Mar 2024 - "The Good, The Bad, and The Ugly: A tale of Physics, Software, and ML", Alex Golub, Advanced Computing and Analysis Techniques in Physics Research (ACAT) Workshop

February, 2024

- 1 Feb 2024 - "Advances in developing deep neural networks for finding primary vertices in proton-proton collisions at the LHC (poster)", Simon Akar, 6th Inter-experiment Machine Learning Workshop

January, 2024

- 31 Jan 2024 - "Towards Differentiable Physics Analysis at the High-Luminosity Large Hadron Collider and Beyond", Matthew Feickert, UC Berkeley Neyman Seminar
- 30 Jan 2024 - "Advances in developing deep neural networks for finding primary vertices in proton-proton collisions at the LHC", Simon Akar, 6th Inter-experiment Machine Learning Workshop
- 10 Jan 2024 - "Three garbage collectors: Java, Python, and Julia", Jim Pivarski, Compute & Accelerator Forum

Papers

Publications by the IRIS-HEP team

- Physics analysis for the HL-LHC: concepts and pipelines in practice with the Analysis Grand Challenge, A. Held, E. Kauffman, O. Shadura and A. Wightman, arXiv 2401.02766 (05 Jan 2024).
- Machine Learning for Columnar High Energy Physics Analysis, E. Kauffman, A. Held and O. Shadura, arXiv 2401.01802 (03 Jan 2024).
- Generalizing mkFit and its Application to HL-LHC, Giuseppe Cerati, Peter Elmer, Patrick Gartung, Leonardo Giannini, Matti Kortelainen, Vyacheslav Krutelyov, Steven Lantz, Mario Masciovecchio, Tres Reid, Allison Reinsvold Hall, Daniel Riley, Matevz Tadel, Emmanouil Vourliotis, Peter Wittich, Avi Yagil (for the CMS Collaboration), arXiv:2312.11728 (2023). (Submitted to CHEP-2023) (18 Dec 2023).
- RenderCore -- a new WebGPU-based rendering engine for ROOT-EVE, Ciril Bohak, Dmytro Kovalskyi, Sergey Linev, Alja Mrak Tadel, Sebastien Strban, Matevz Tadel, Avi Yagil, arXiv.2312.11729 (2023). (Submitted to CHEP-2023) (18 Dec 2023).
- High Pileup Particle Tracking with Object Condensation, Kilian Lieret, Gage DeZoort, Devdoot Chatterjee, Jian Park, et.al. High Pileup Particle Tracking with Object Condensation. Submitted to 8th International Connecting The Dots Workshop (Toulouse 2023) (06 Dec 2023).
- Train To Sustain, Sudhir Malik, Kilian Lieret, Peter Elmer, Michel Hernandez Villanueva, and Stefan Roiser (15 Nov 2023).
- Training and Onboarding initiatives in High Energy Physics experiments, S. Hageboeck, A. Reinsvold Hall, N. Skidmore, G. A. Stewart, G. Benelli, B. Carlson, C. David, J. Davies, W. Deconinck, D. DeMuth Jr., P. Elmer, R. B. Garg, K. Lieret, V. Lukashenko, S. Malik, A. Morris, H. Schellman, J. Veatch, M. Hernandez Villanueva (23 Oct 2023).
- Awkward Just-In-Time (JIT) Compilation: A Developer's Experience, Ianna Osborne, Jim Pivarski, Ioana Ifrim, Angus Hollands and Henry Schreiner, arXiv:2310.01461 [cs.PL] (Submitted to CHEP 2023) (03 Oct 2023).
- Bayesian Methodologies with pyhf, M. Feickert, L. Heinrich and M. Horstmann, arXiv 2309.17005 (29 Sep 2023) [1 citation].
- An Object Condensation Pipeline for Charged Particle Tracking, Kilian Lieret and Gage DeZoort. An Object Condensation Pipeline for Charged Particle Tracking. Submitted to 8th International Connecting The Dots Workshop (Toulouse 2023) (28 Sep 2023).
- Software Citation in HEP: Current State and Recommendations for the Future, M. Feickert, D. Katz, M. Neubauer, E. Sexton-Kennedy and G. Stewart, arXiv 2309.14571 (25 Sep 2023).
- Software Training Outreach In HEP , Sudhir Malik, Danelix Cordero, Peter Elmer, Adam LaMee, and Ken Cecire (24 Sep 2023).
- Advances in developing deep neural networks for finding primary vertices in proton-proton collisions at the LHC, S. Akar, M. Elashri, R. Garg, E. Kauffman, M. Peters, H. Schreiner, M. Sokoloff, W. Tepe and L. Tompkins, arXiv 2309.12417 (21 Sep 2023).
- Analyzing Transatlantic Network Traffic over Scientific Data Caches, Ziyue Deng, Alex Sim, Kesheng Wu, Chin Guok, Inder Monga, Fabio Andrijauskas, Frank Wuerthwein, and Derek Weitzel. Analyzing Transatlantic Network Traffic over Scientific Data Caches. in ACM 6th International Workshop on System and Network Telemetry and Analytics (SNTA'23) (20 Jun 2023).
- Fast And Automatic Floating Point Error Analysis With CHEF-FP, Garima Singh, Baidyanath Kundu, Harshitha Menon, Alexander Penev, David J. Lange and Vassil Vassilev, arXiv:2304.06441 [math.NA] (Submitted to IEEE International Parallel and Distributed Processing Symposium 2023) (13 Apr 2023).

Conclusions

- With ~30 FTE's we have a lot of on-going projects!
  - And many places to collaborate
  - Nothing to stop new efforts from starting
- Lots of projects with possible synergies in Analysis Systems and DOMA
- Not mentioned: Innovative Algorithms
  - Work on tracking (ACTS for ATLAS)
  - With some intersection with ML
- Talk to me and anyone else who is part of IRIS-HEP: we can direct you to the appropriate contact points
- This is v2 of IRIS-HEP and we are focused on delivery more than R&D