

Analysis Facilities intro

A. Forti

GridPP51

26 March 2024



Evolution of AF definition

- Analysis facilities have often been considered as something separate from the current grid infrastructure, mainly due to few aspects:
 - Large scale interactivity for fast analysis cycles
 - Integration of cloud native technologies
 - Concentrating on python based parallelized workflows on reduced data formats (nanoAOD/PHYSLITE)
 - Assumption that data formats small enough to be local to the AF resources
- Reality is more complicated and the community agreed the AF have to be integrated with the existing infrastructure and preferably co-located to grid sites.



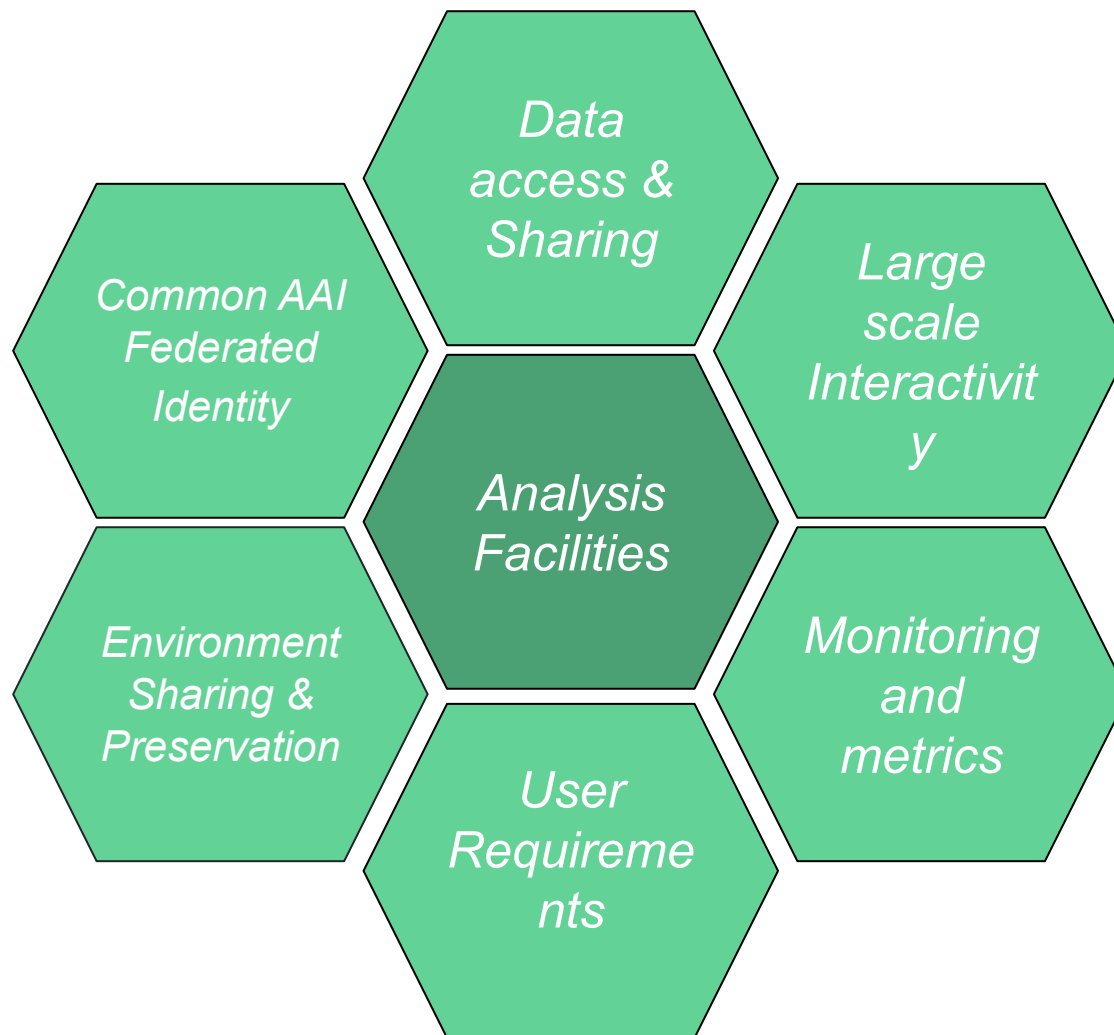
Analysis Facilities

Infrastructure and services that provide integrated data, software and computational resources to execute one or more elements of an analysis workflow. These resources are shared among members of a virtual organization and supported by that organization.

A set of R&Ds to improve

Scalability
Ease of use
Data Access
Interoperability
Collaboration

For the end user analysis at the
HL-LHC scale



Analysis Facilities White Paper

- Contains the review of the R&D on analysis.
 - Since last year [pre-CHEP workshop](#) reviewed and incorporated >100 comments from all parties
- Now in [finalised stage](#)
 - Only endorsements
- To be published on arXiv on the 2nd of April.
- Next stage will start at the [WLCG/HSF workshop in May](#)

1

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], Andr s Pacheco Pages [11], Antonio Delgado Peris [12], Nick Smith [13], Jonas Eschle [14], Liz Sexton-Kennedy [13], Evangelos Kourlitis [3], Alexander Held [15], Jos  Flix Molina [12], Antonio Perez-Calero Yzquierdo [12], Burt Holzman [13], Niladri B hoo [16], Kevin Pedro [13], Daniele Spiga [1], Jamie Gooding [17], Giordon Stark [18], Clemens Lange [19], Piergiulio Lenzi [1], Thomas Kuhr [27], 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], Jos  Hern ndez [12], Brian Bockelman [25], Eric Lancon [9], James Frost [26], Andrea Sciaba [6], Markus Schulz [6], Caterina Doglioni [2], Guenter Duckeck [27], Tibor S mko [6], Luke Kreczko [28], Farid Ould-Saada [29], Eirik Gramstad [29], James Catmore [29], Matthew Feickert [15], Iason Krommydas [30], Shawn McKee [31], Tomas Lind n [32], Jim Pivarski [33], Peter Elmer [33]

b editor, 1 INFN, 2 University of Manchester, 3 Technische Universit t M nchen, 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 Universit t 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, 26 University of Oxford, 27 Ludwigs-Maximilians Universit t M nchen, 28 University of Bristol, 29 University of Oslo, 30 Rice University, 31 University of Michigan, 32 Helsinki Institute of Physics, 33 Princeton University

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



LHCC/WLCG/HSF

The Analysis Facilities is a community effort and is widespread and the [LHCC has requested experiments](#) to engage with the R&D and give a direction.

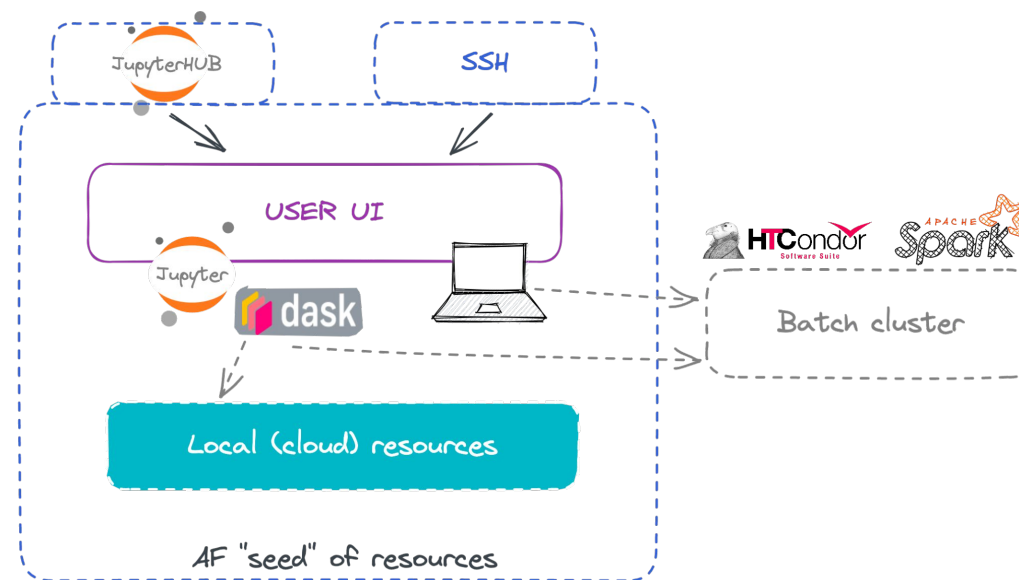
WLCG has now incorporated Analysis Facilities in the [strategy document](#)
[WLCG/HSF workshop](#) sessions on AF and a lot of strategy discussion.




Scaled up interactivity

Dask and GPUs

- Accepted view is to have a seed of interactive resources and offload to a batch system/grid after the initial development from the same interface
 - Included dask jobs and jobs accessing GPUs



- Luke will present the tests done by Swift-HEP to try to apply this model using Dirac



GridPP/SWIFT-HEP joint meeting

Analysis Facilities

An overview

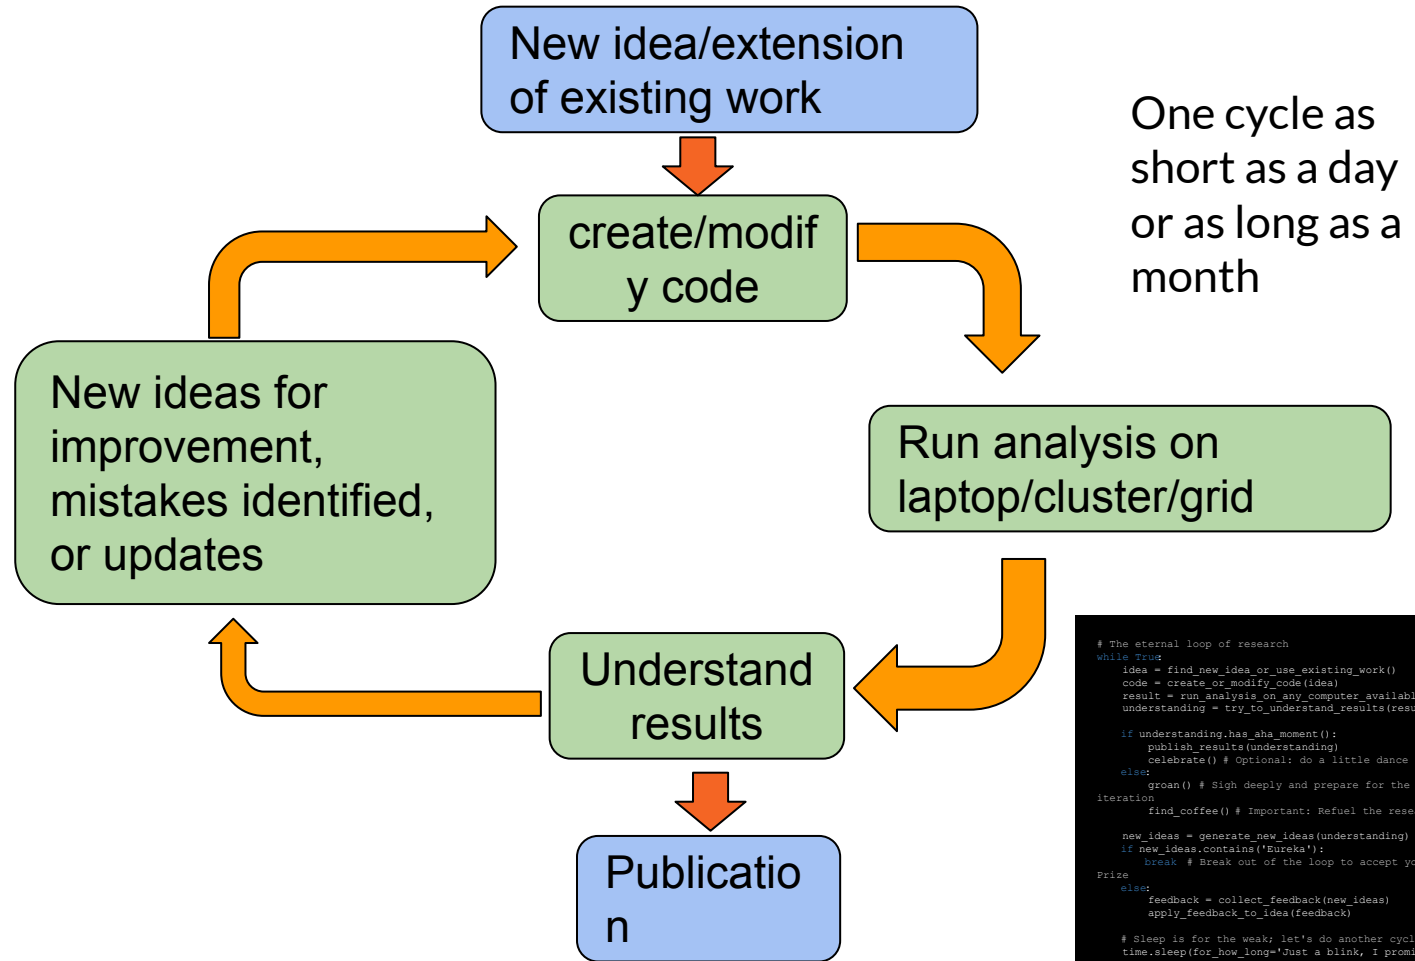
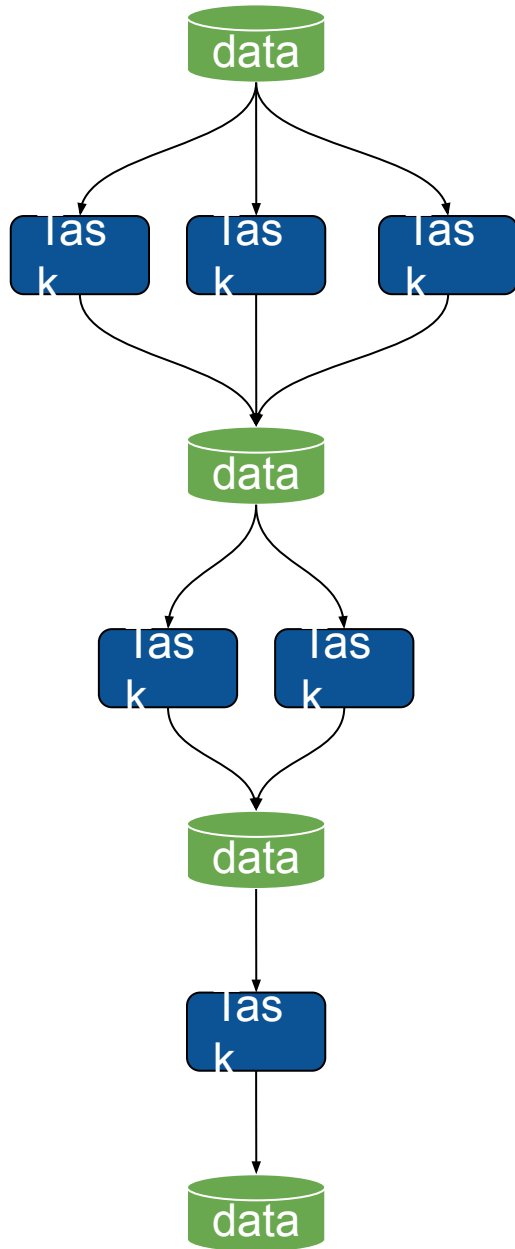
Luke Kreczko



Anatomy of an analysis

workflow

The cycle (oversimplified)



```
# The eternal loop of research
while True:
    idea = find_new_idea_or_use_existing_work()
    code = create_or_modify_code(idea)
    result = run_analysis_on_any_computer_available(code)
    understanding = try_to_understand_results(result)


    if understanding.has_aha_moment():
        publish_results(understanding)
        celebrate() # Optional: do a little dance
    else:
        groan() # Sigh deeply and prepare for the next iteration
        find_coffee() # Important: Refuel the researcher

    new_ideas = generate_new_ideas(understanding)
    if new_ideas.contains('Eureka'):
        break # Break out of the loop to accept your Nobel Prize
    else:
        feedback = collect_feedback(new_ideas)
        apply_feedback_to_idea(feedback)

# Sleep is for the weak; let's do another cycle!
time.sleep(for_how_long='Just a blink, I promise')

# Note: This loop does not account for weekends or holidays
# which are often mythical in research
```


SWIFT-HEP context



Use existing resources (GridPP) to construct a (virtual) analysis facility

- Dask as user interface
- DIRAC for creating and scaling the cluster
- Rucio for data replication

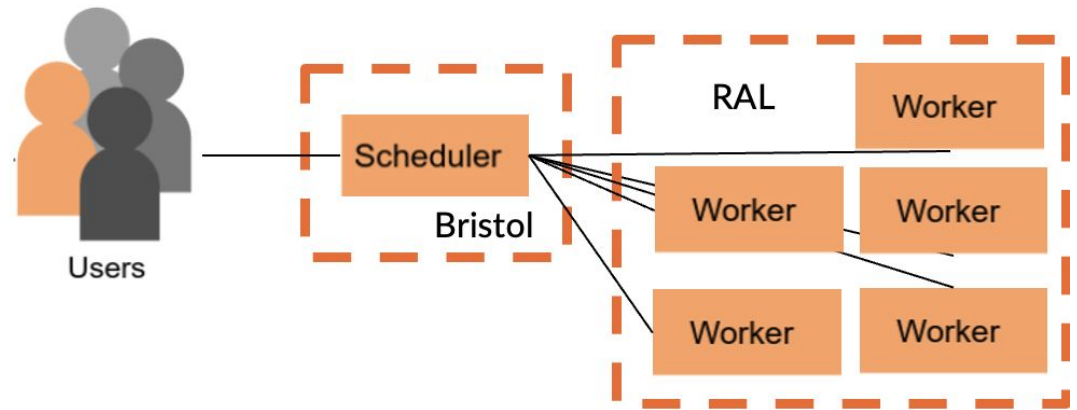
Simplify grid resource access:

- Submit via Jupyter notebooks
- Use Dask tools for interactive progress map
- Without voms-proxy-init?

Tests with RAL PPD

RAL has nodes which have open ports. Will have to discuss security settings to run the prototype on multiple sites

setting	number of files	total size	number of events
1	9	22.9 GB	10,455,719
2	18	42.8 GB	19,497,435
5	43	105 GB	47,996,231
10	79	200 GB	90,546,458
20	140	359 GB	163,123,242
50	255	631 GB	297,247,463
100	395	960 GB	470,397,795
200	595	1.40 TB	705,273,291
-1	787	1.78 TB	940,160,174

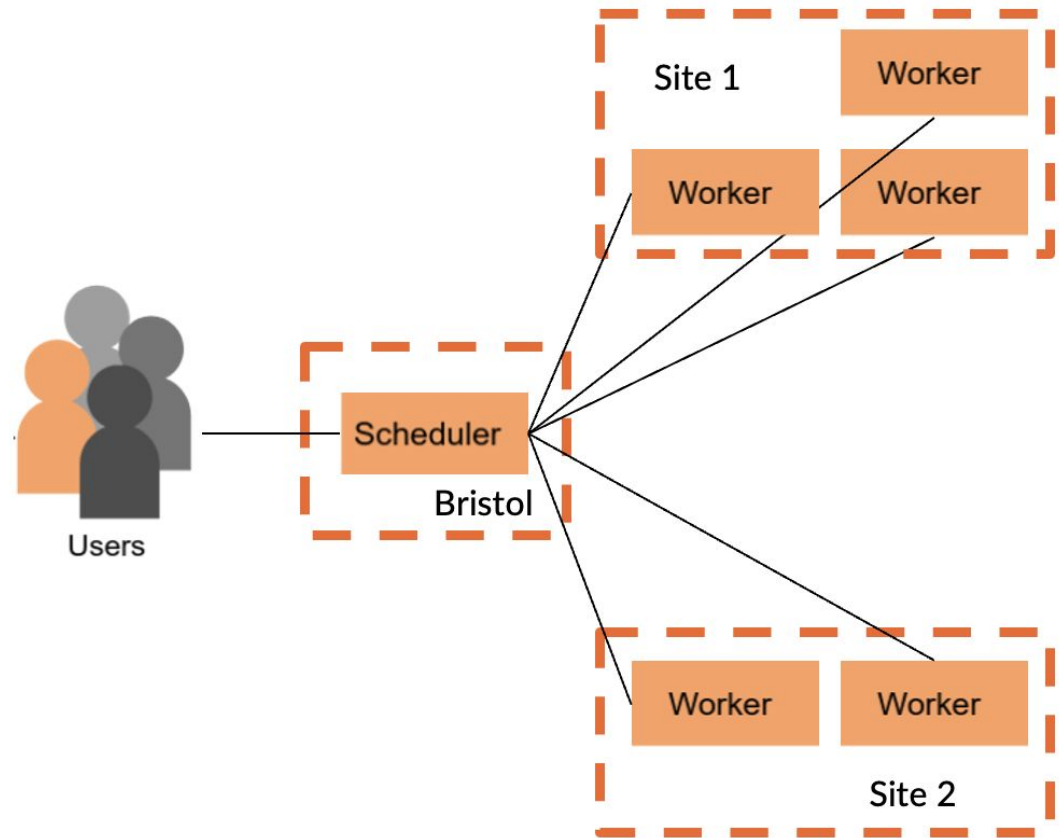


Going beyond a single site

Use the Analysis Grand Challenge as a test case (CMS data)

Workers and Scheduler needed to be able to talk to each other -> firewalls stop prevent this traffic

Science DMZ (all GridPP) would be great; Software alternatives (e.g. [Zenith](#)) exist



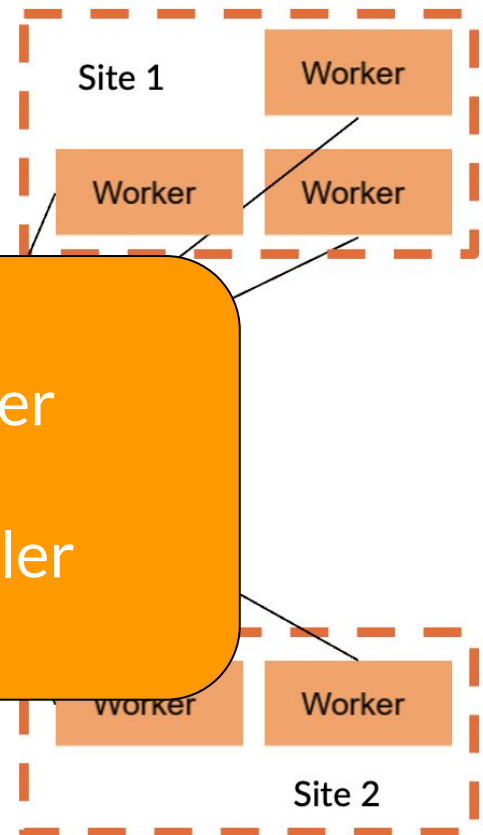
Going beyond a single site

Use the Analysis Grand Challenge as a test case (CMS data)

Workers and
be able to talk
firewalls stop


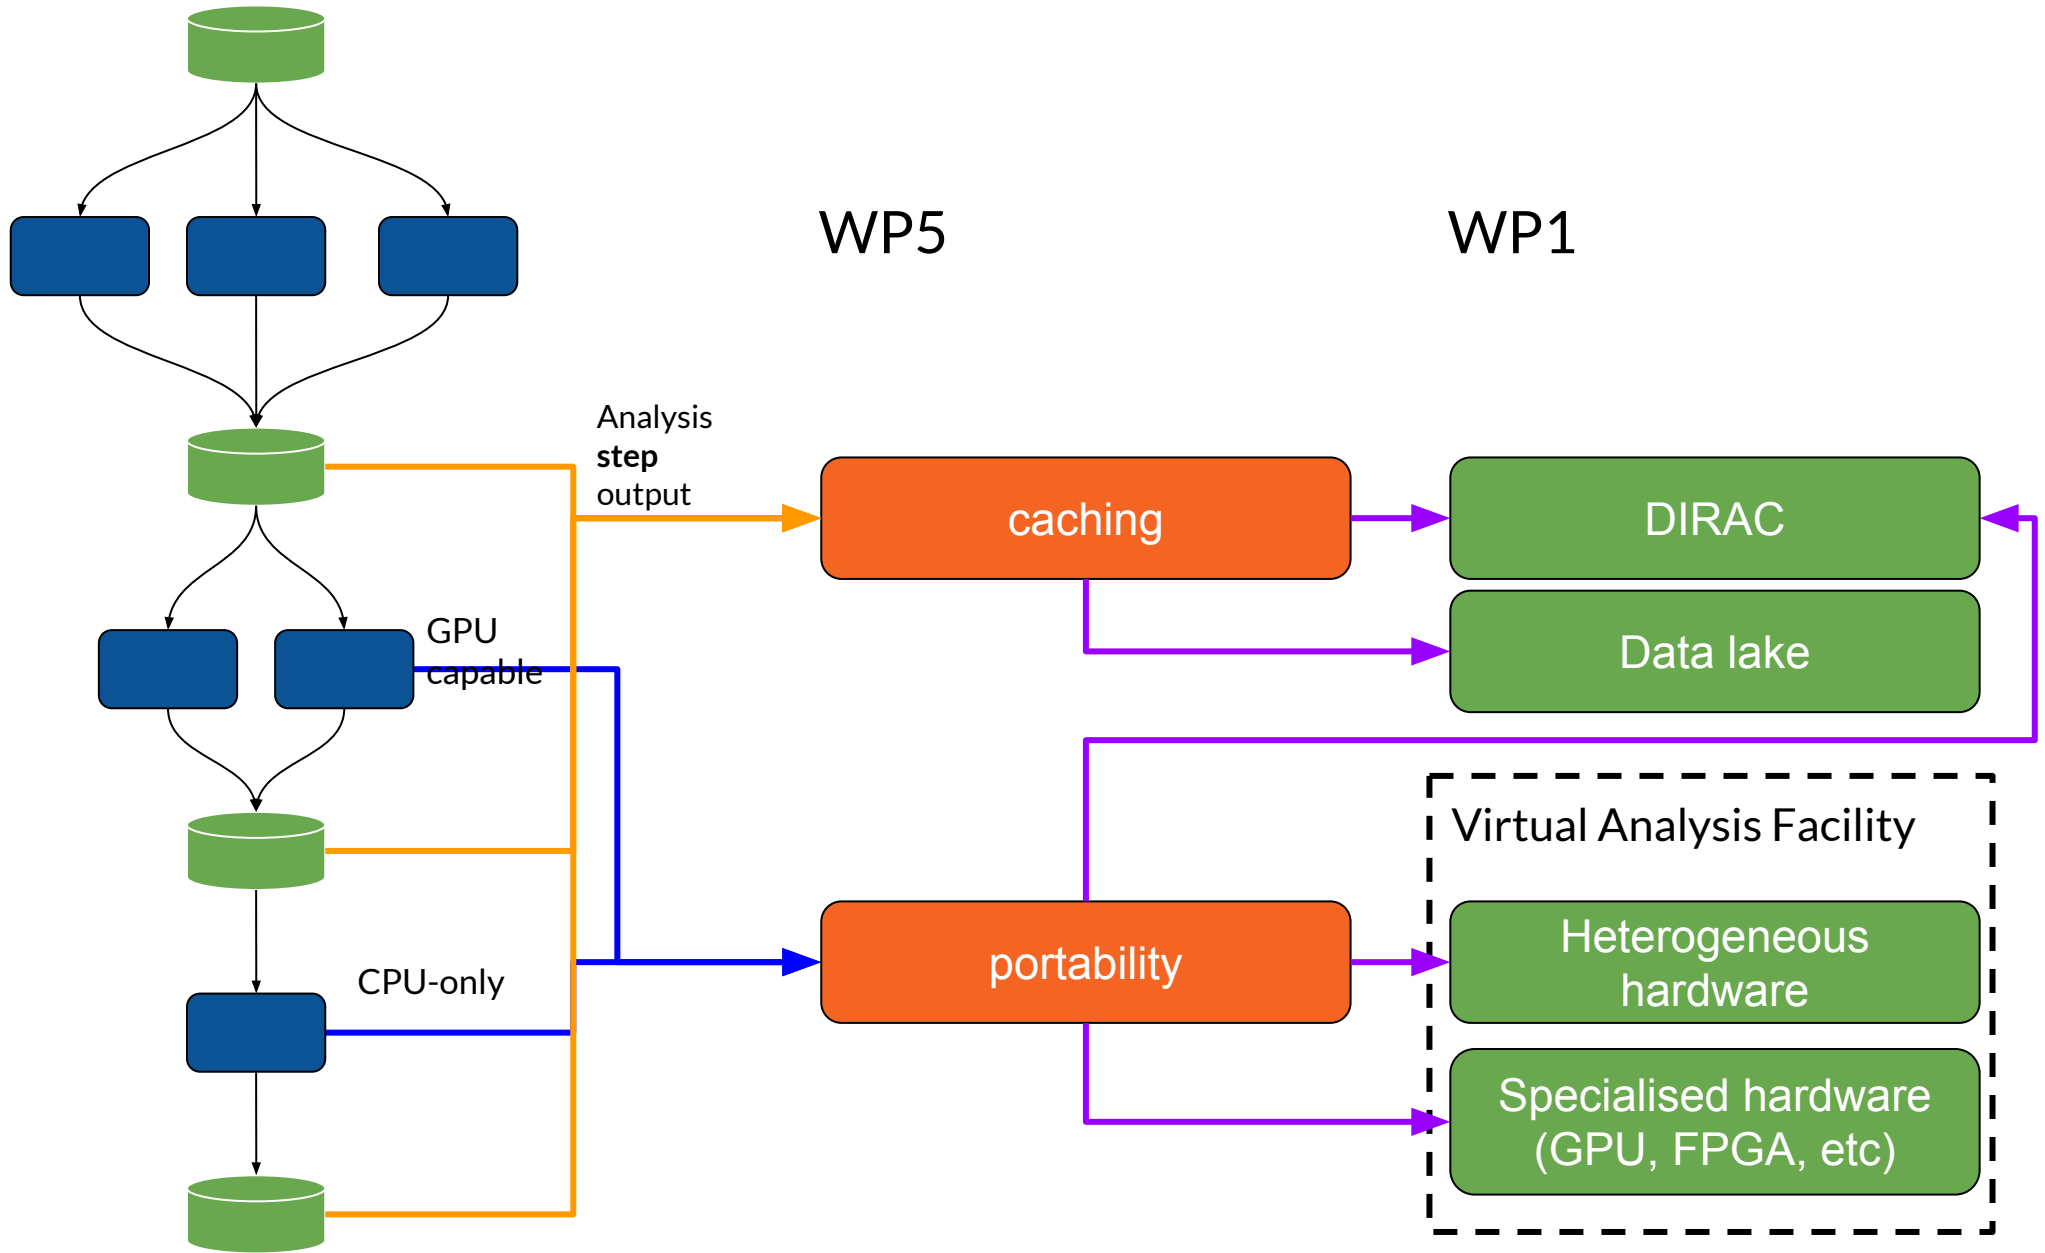
Science DMZ (all GridPP) would be great

Note: Efficiency of such a split is under investigation.
All data is transferred back to scheduler




SWIFT-HEP phase 1/1.5



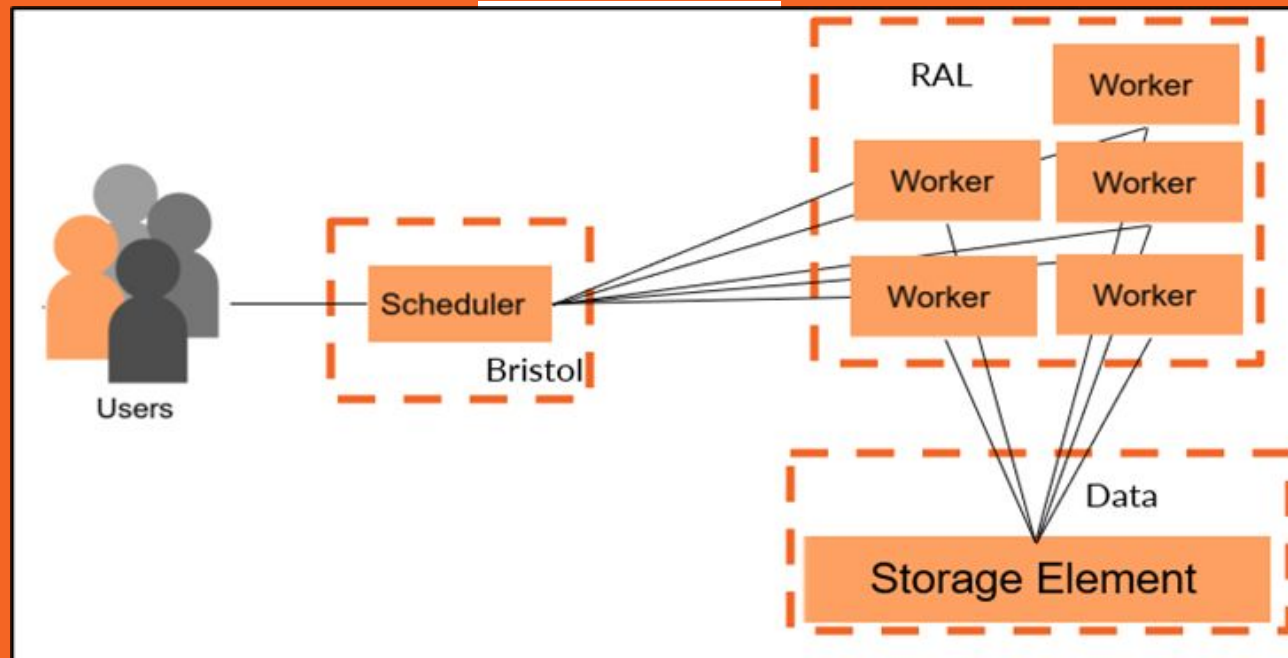
Analysis workflow



Real analysis a bit more complicated. See [“Fine-Grained HEP Analysis Task Graph Optimization with Coffea and Dask”](#) For a more detailed view



For more details, see [Sam's talk tomorrow afternoon](#)



Summary

SWIFT-HEP (WP1+5) is working on a virtual analysis facility

- Use FOSS tools (e.g Dask) to simplify access to the grid
- Understand the benefits and downsides of such a system
- Build with other developments in mind (e.g. IRIS-HEP)

SWIFT-HEP phase 2:

AI-optimized

