Analysis Facilities

ATLAS+CMS Computing BMBF Annual Meeting

Lukas Heinrich

Intro

LHC Computing is a huge success. From a very high level it democratized access to CPU compute + storage

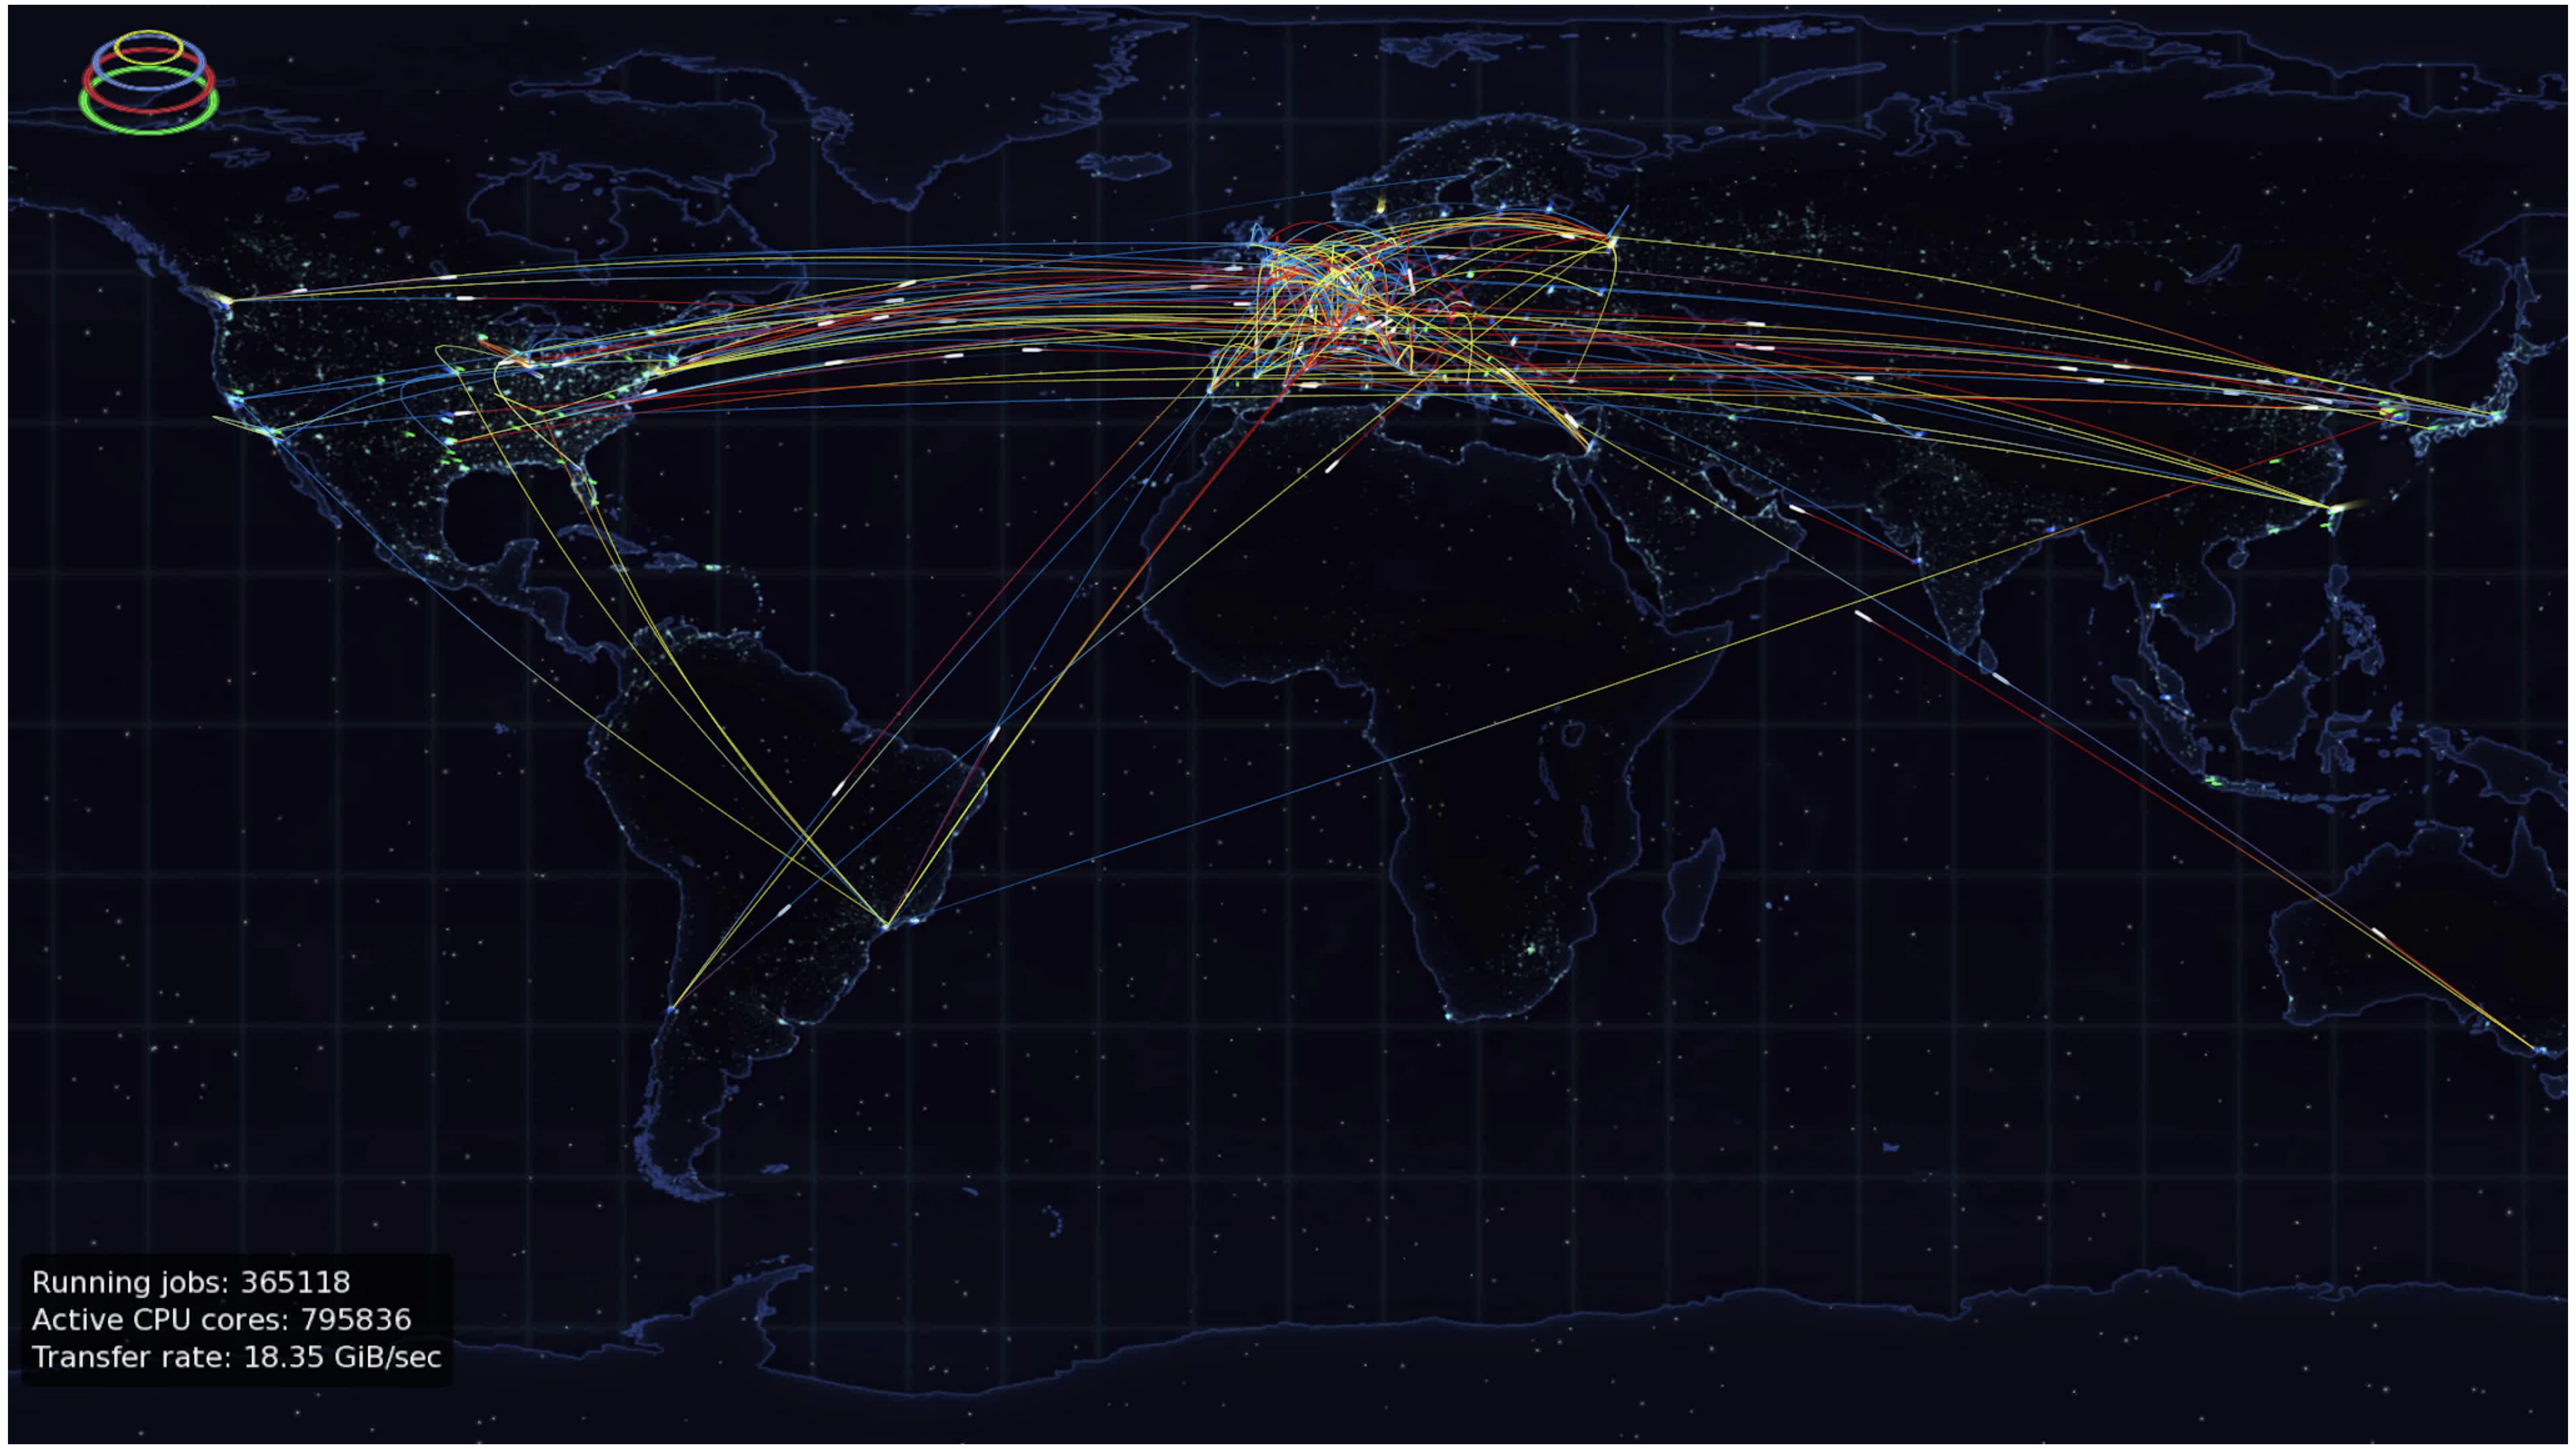

Intro

All this time we’ve always had “Analysis Facilities” of some sort. Infrastructure that enables users to carry out their physics

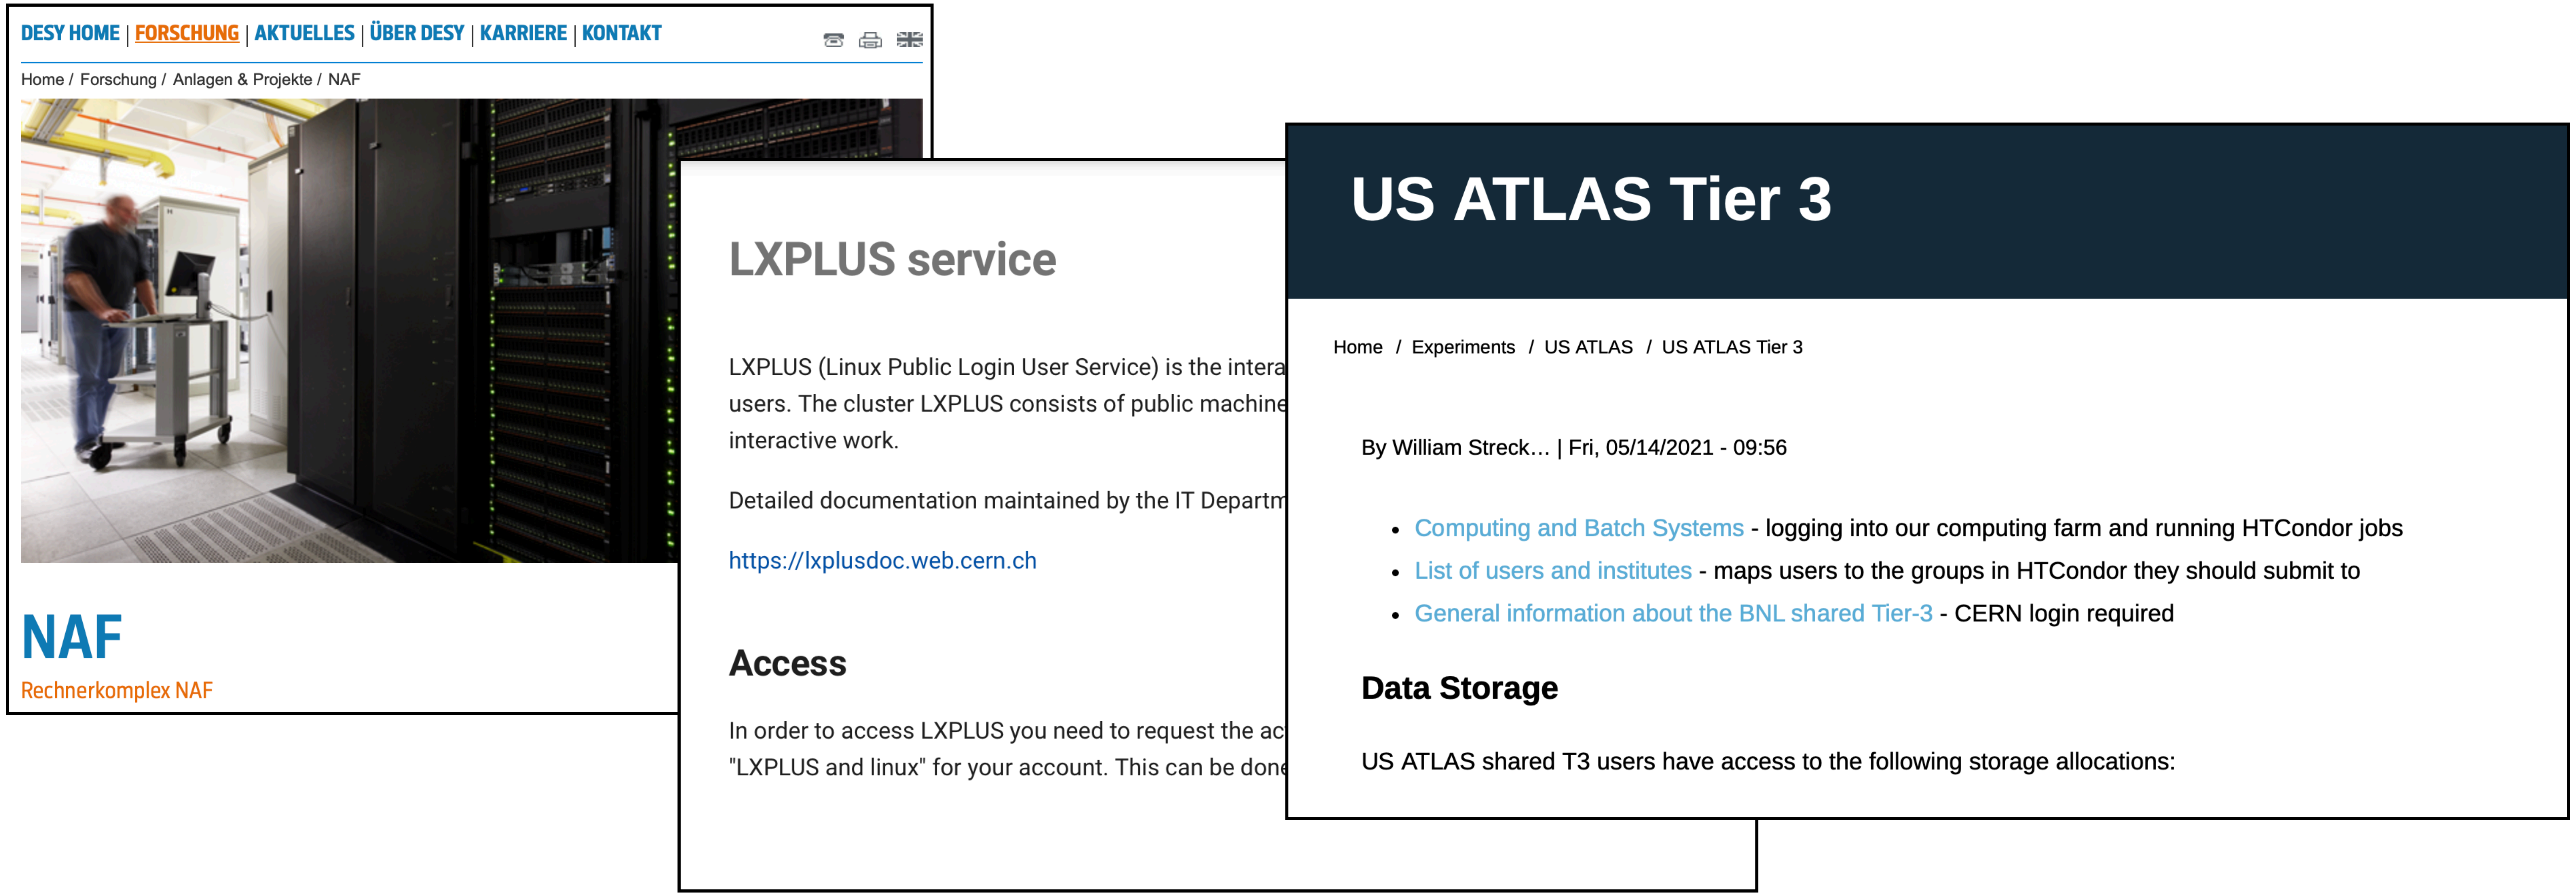

Why are people talking about AFs and what do they mean?

Intro

The model we settled on is very robust + scalable…

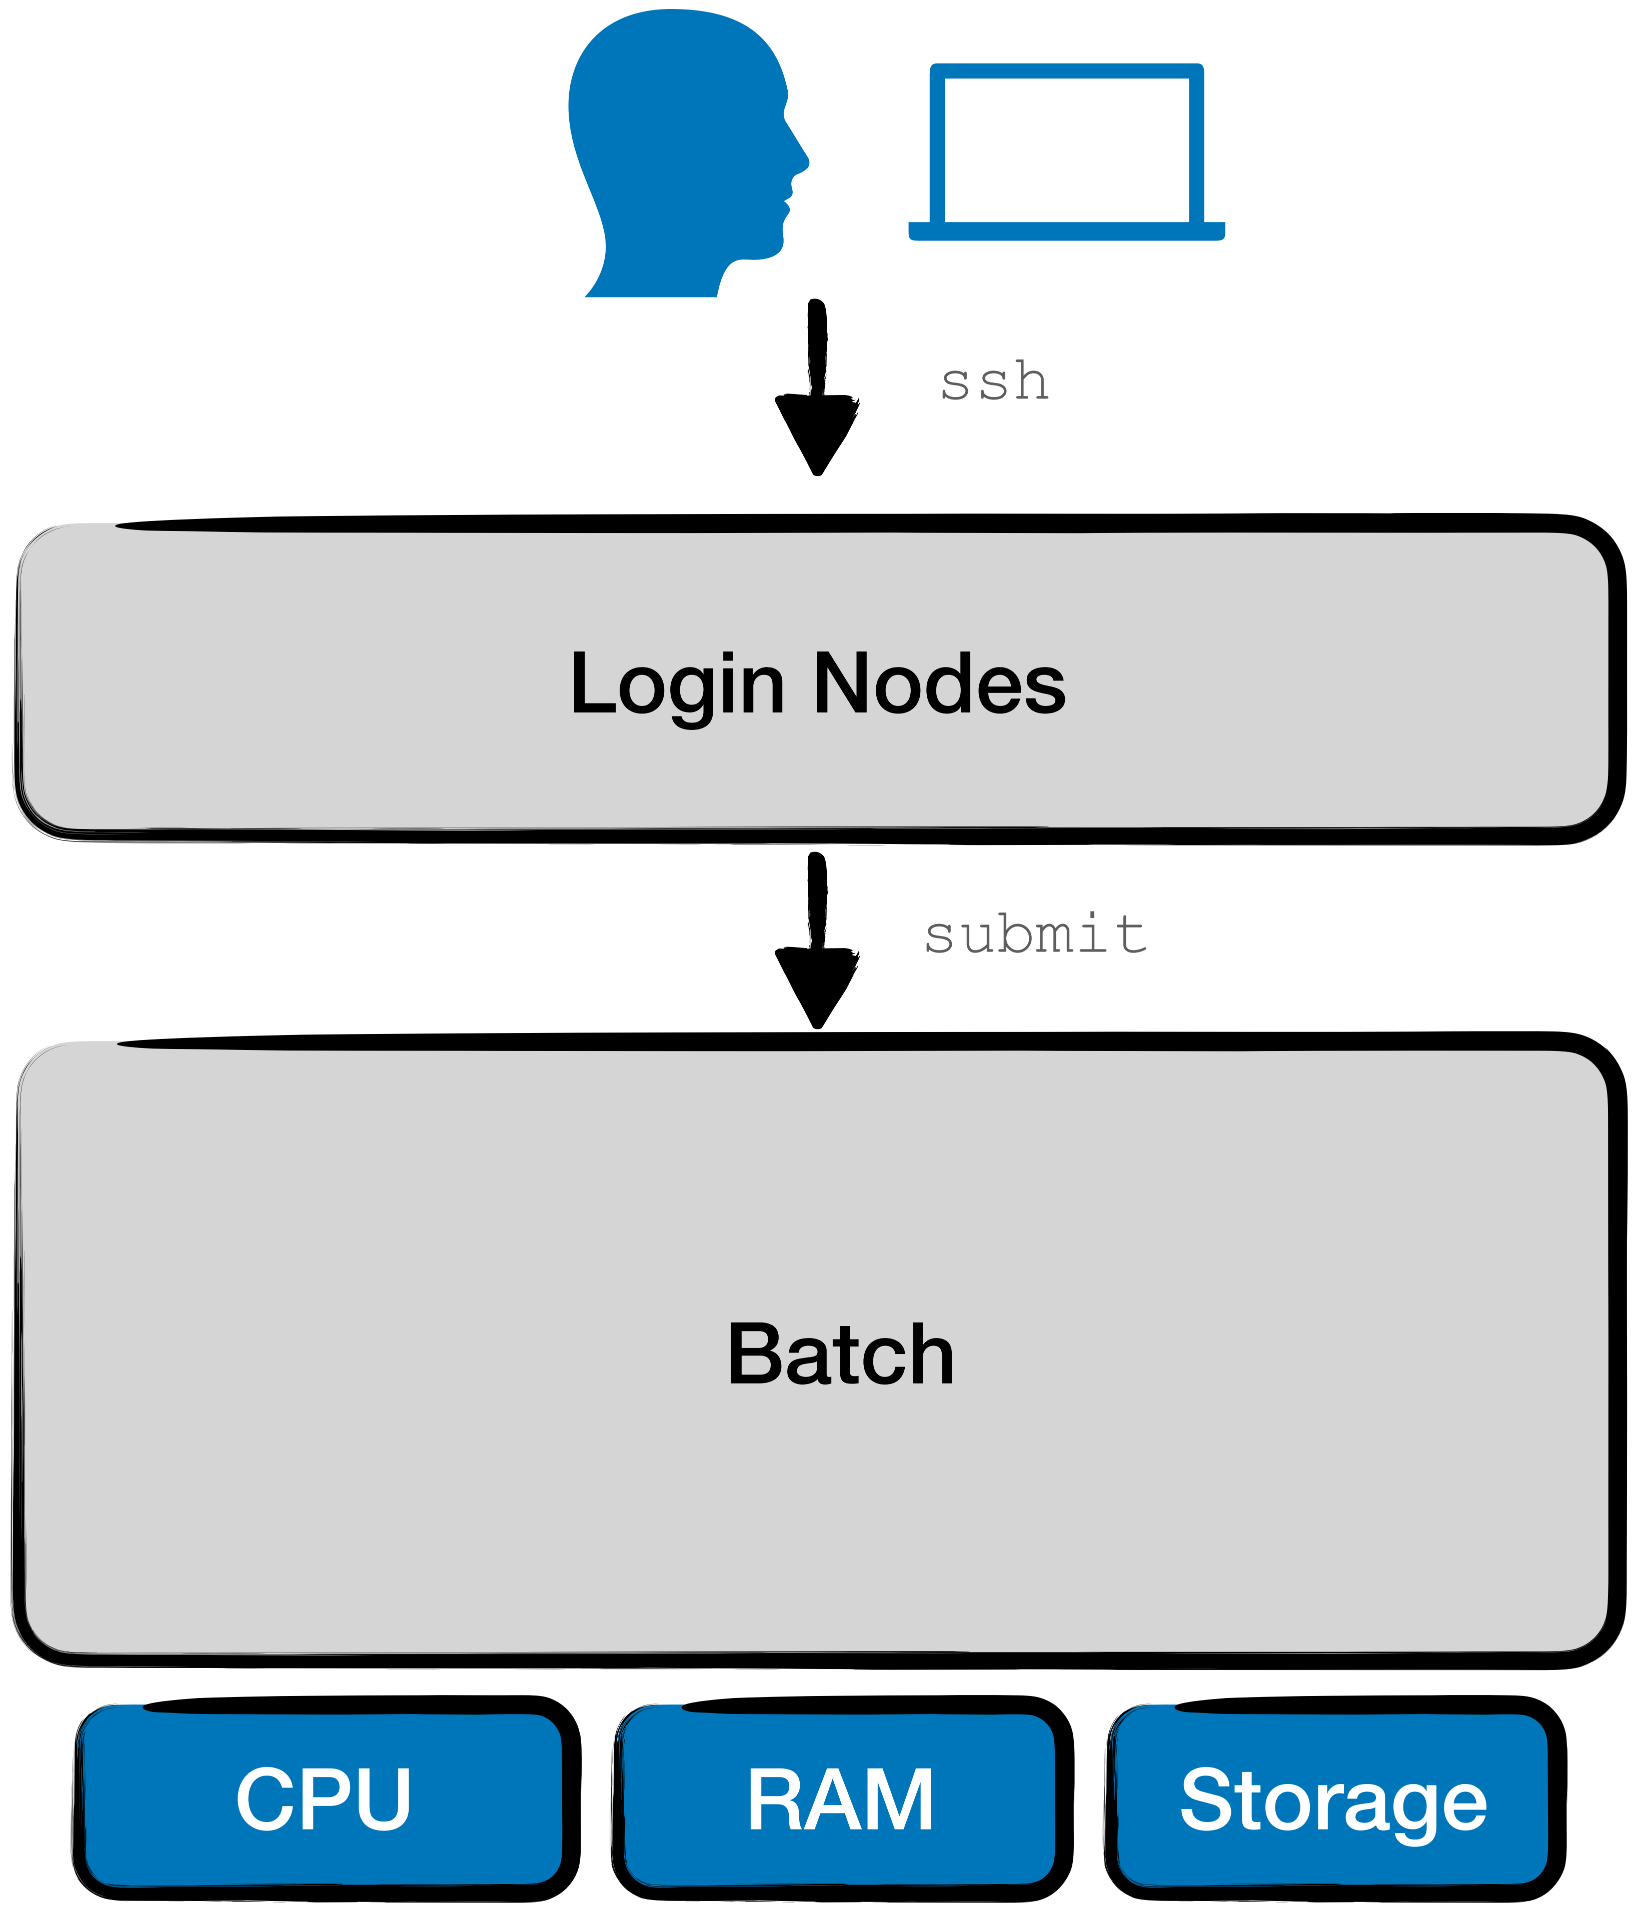

Reminder

… but good to remember that this wasn't always what we envisioned. There were always ambitions for "more"

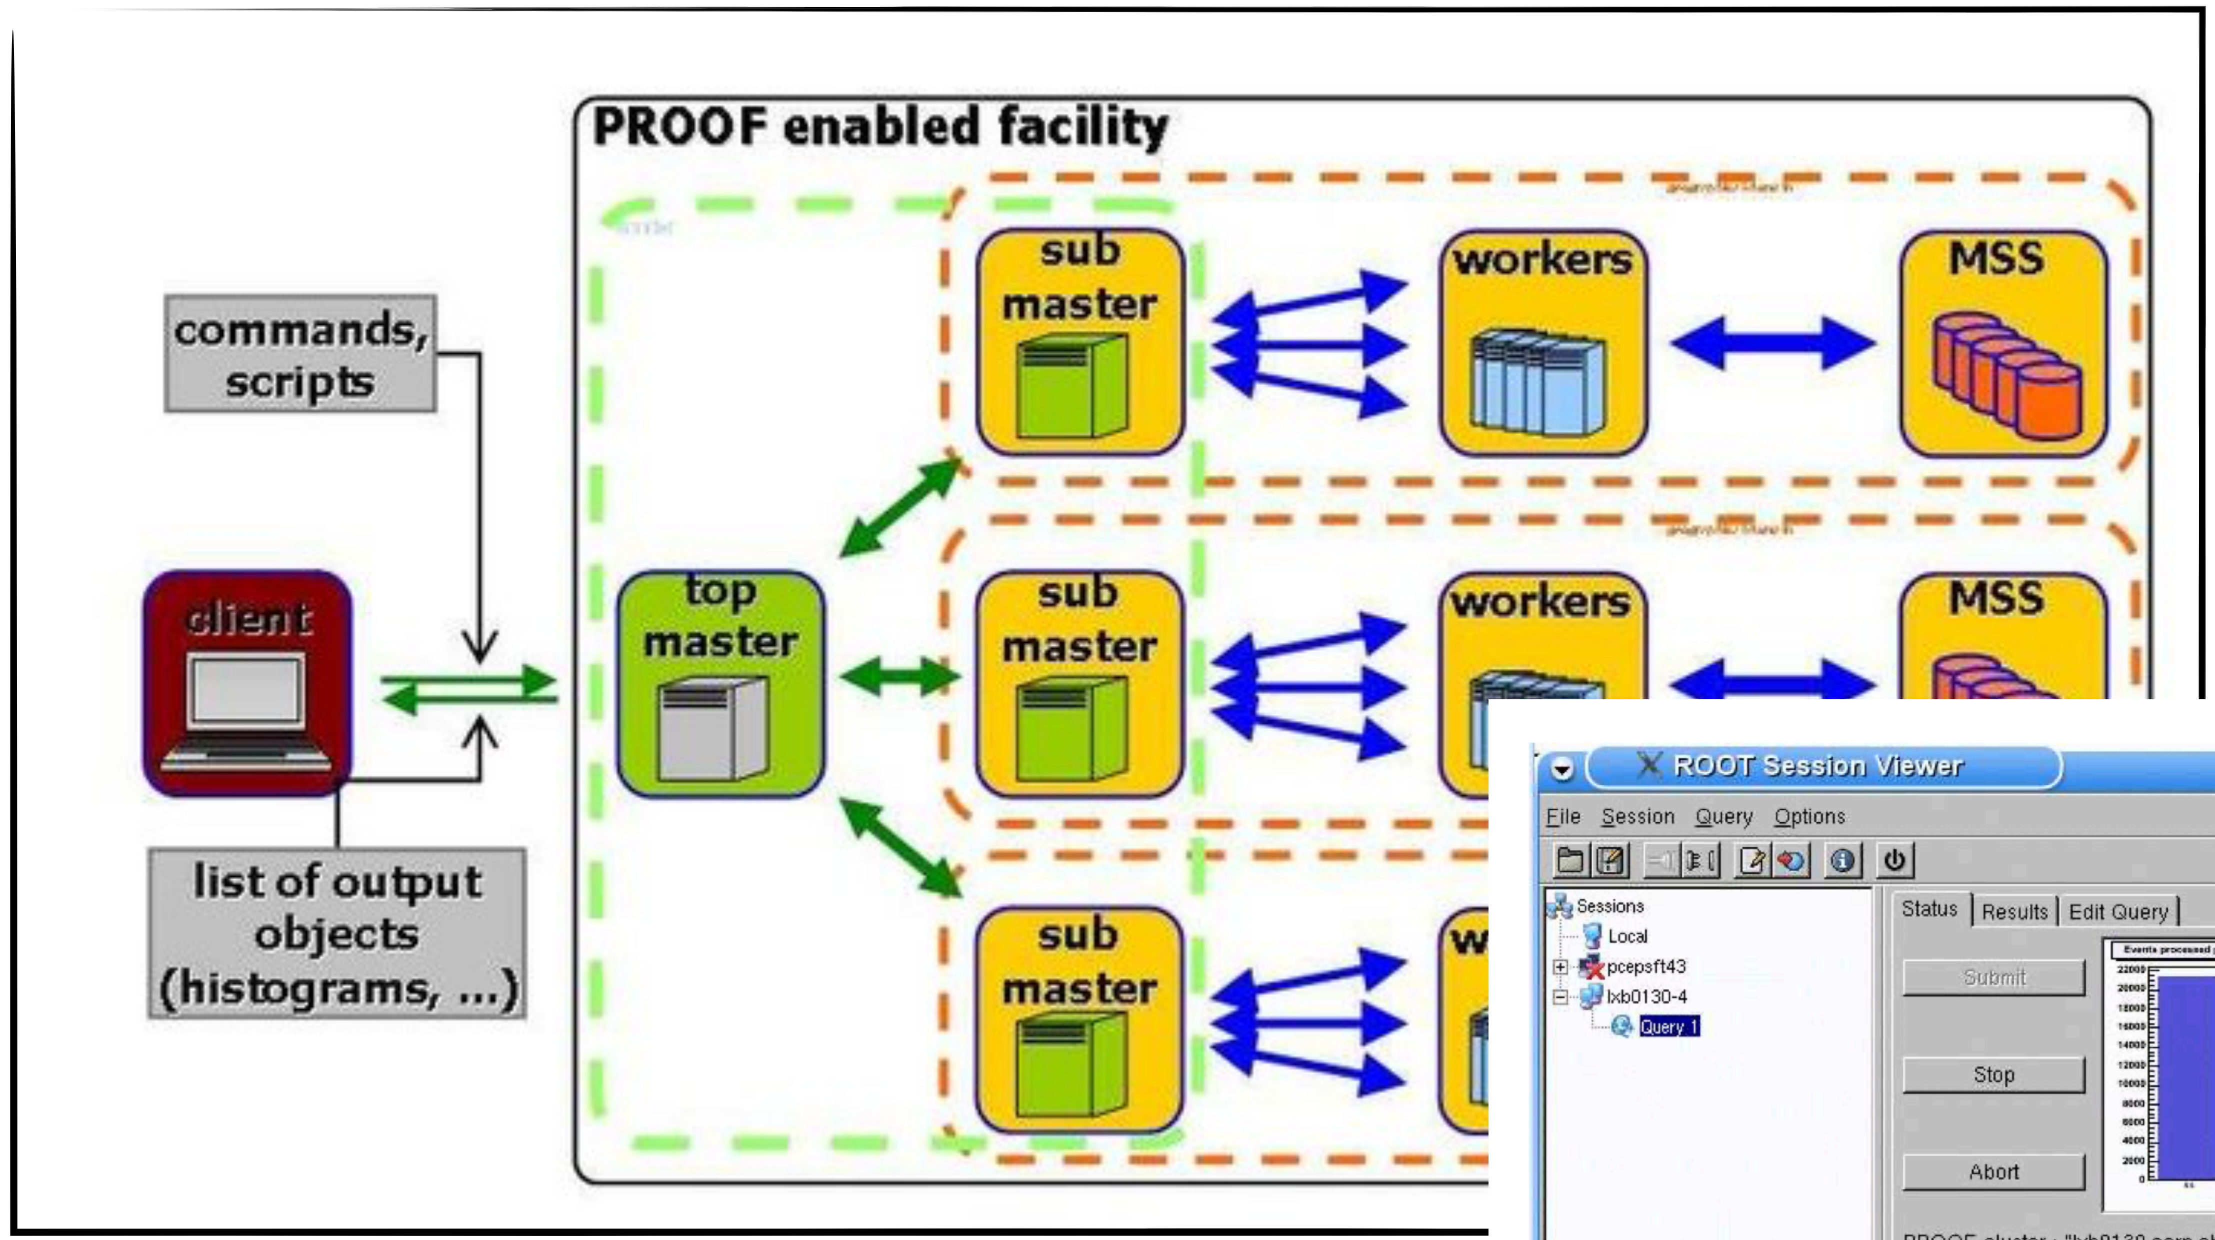

2003

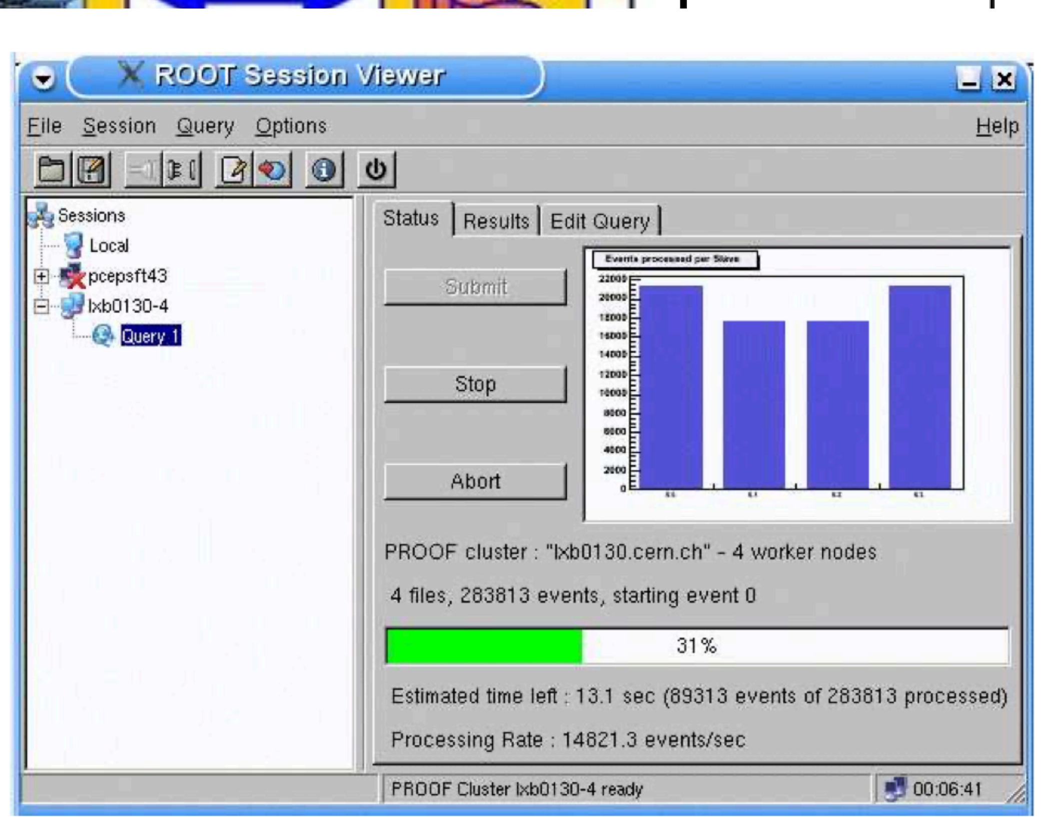

2006 CHEP

ure 3: The session viewer can be used to control multiple PROOF sessio

2005 CMS TDR

4.8.7 PROOF and interactive analysis

The CMS Event Data Model (EDM) is currently evolving in the direction of an event persistence that allows ROOT to be used directly for interactive analysis of standard CMS event data.

This opens up the possibility to use the Parallel ROOT Facility (PROOF) as an integral part of the CMS distributed computing system. At present, we expect that some Tier-2 centres a
for their

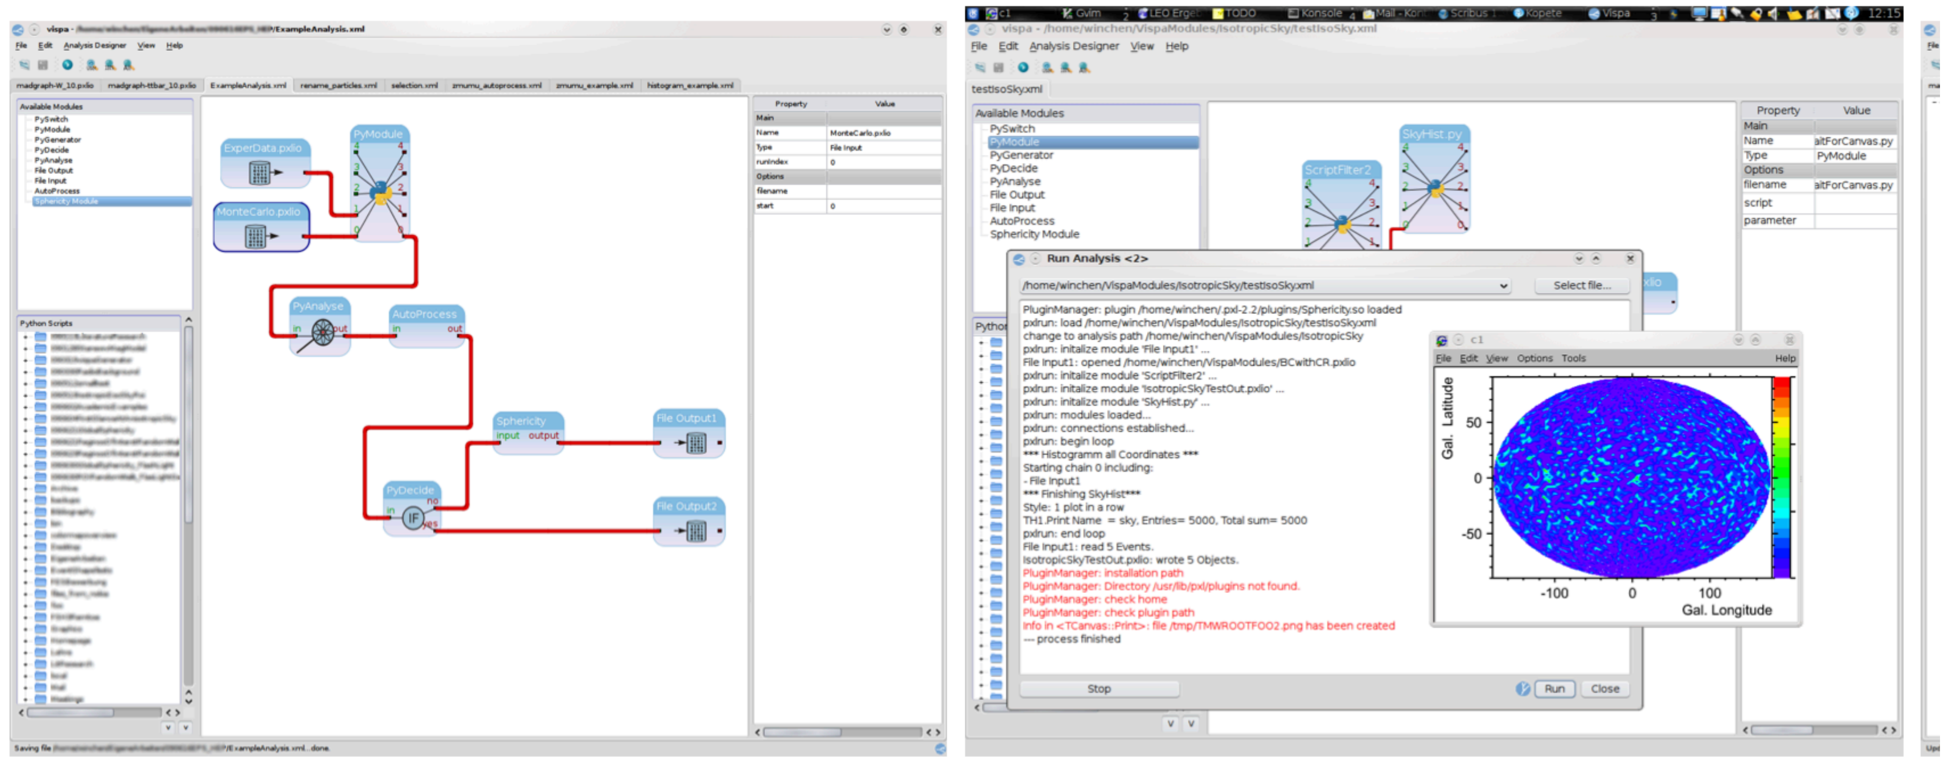

2008 EPS

not necessarily bad ideas, HEP = first mover's disadvantage

The inside world:

HL-LHC: reason to revisit our computing model more generally

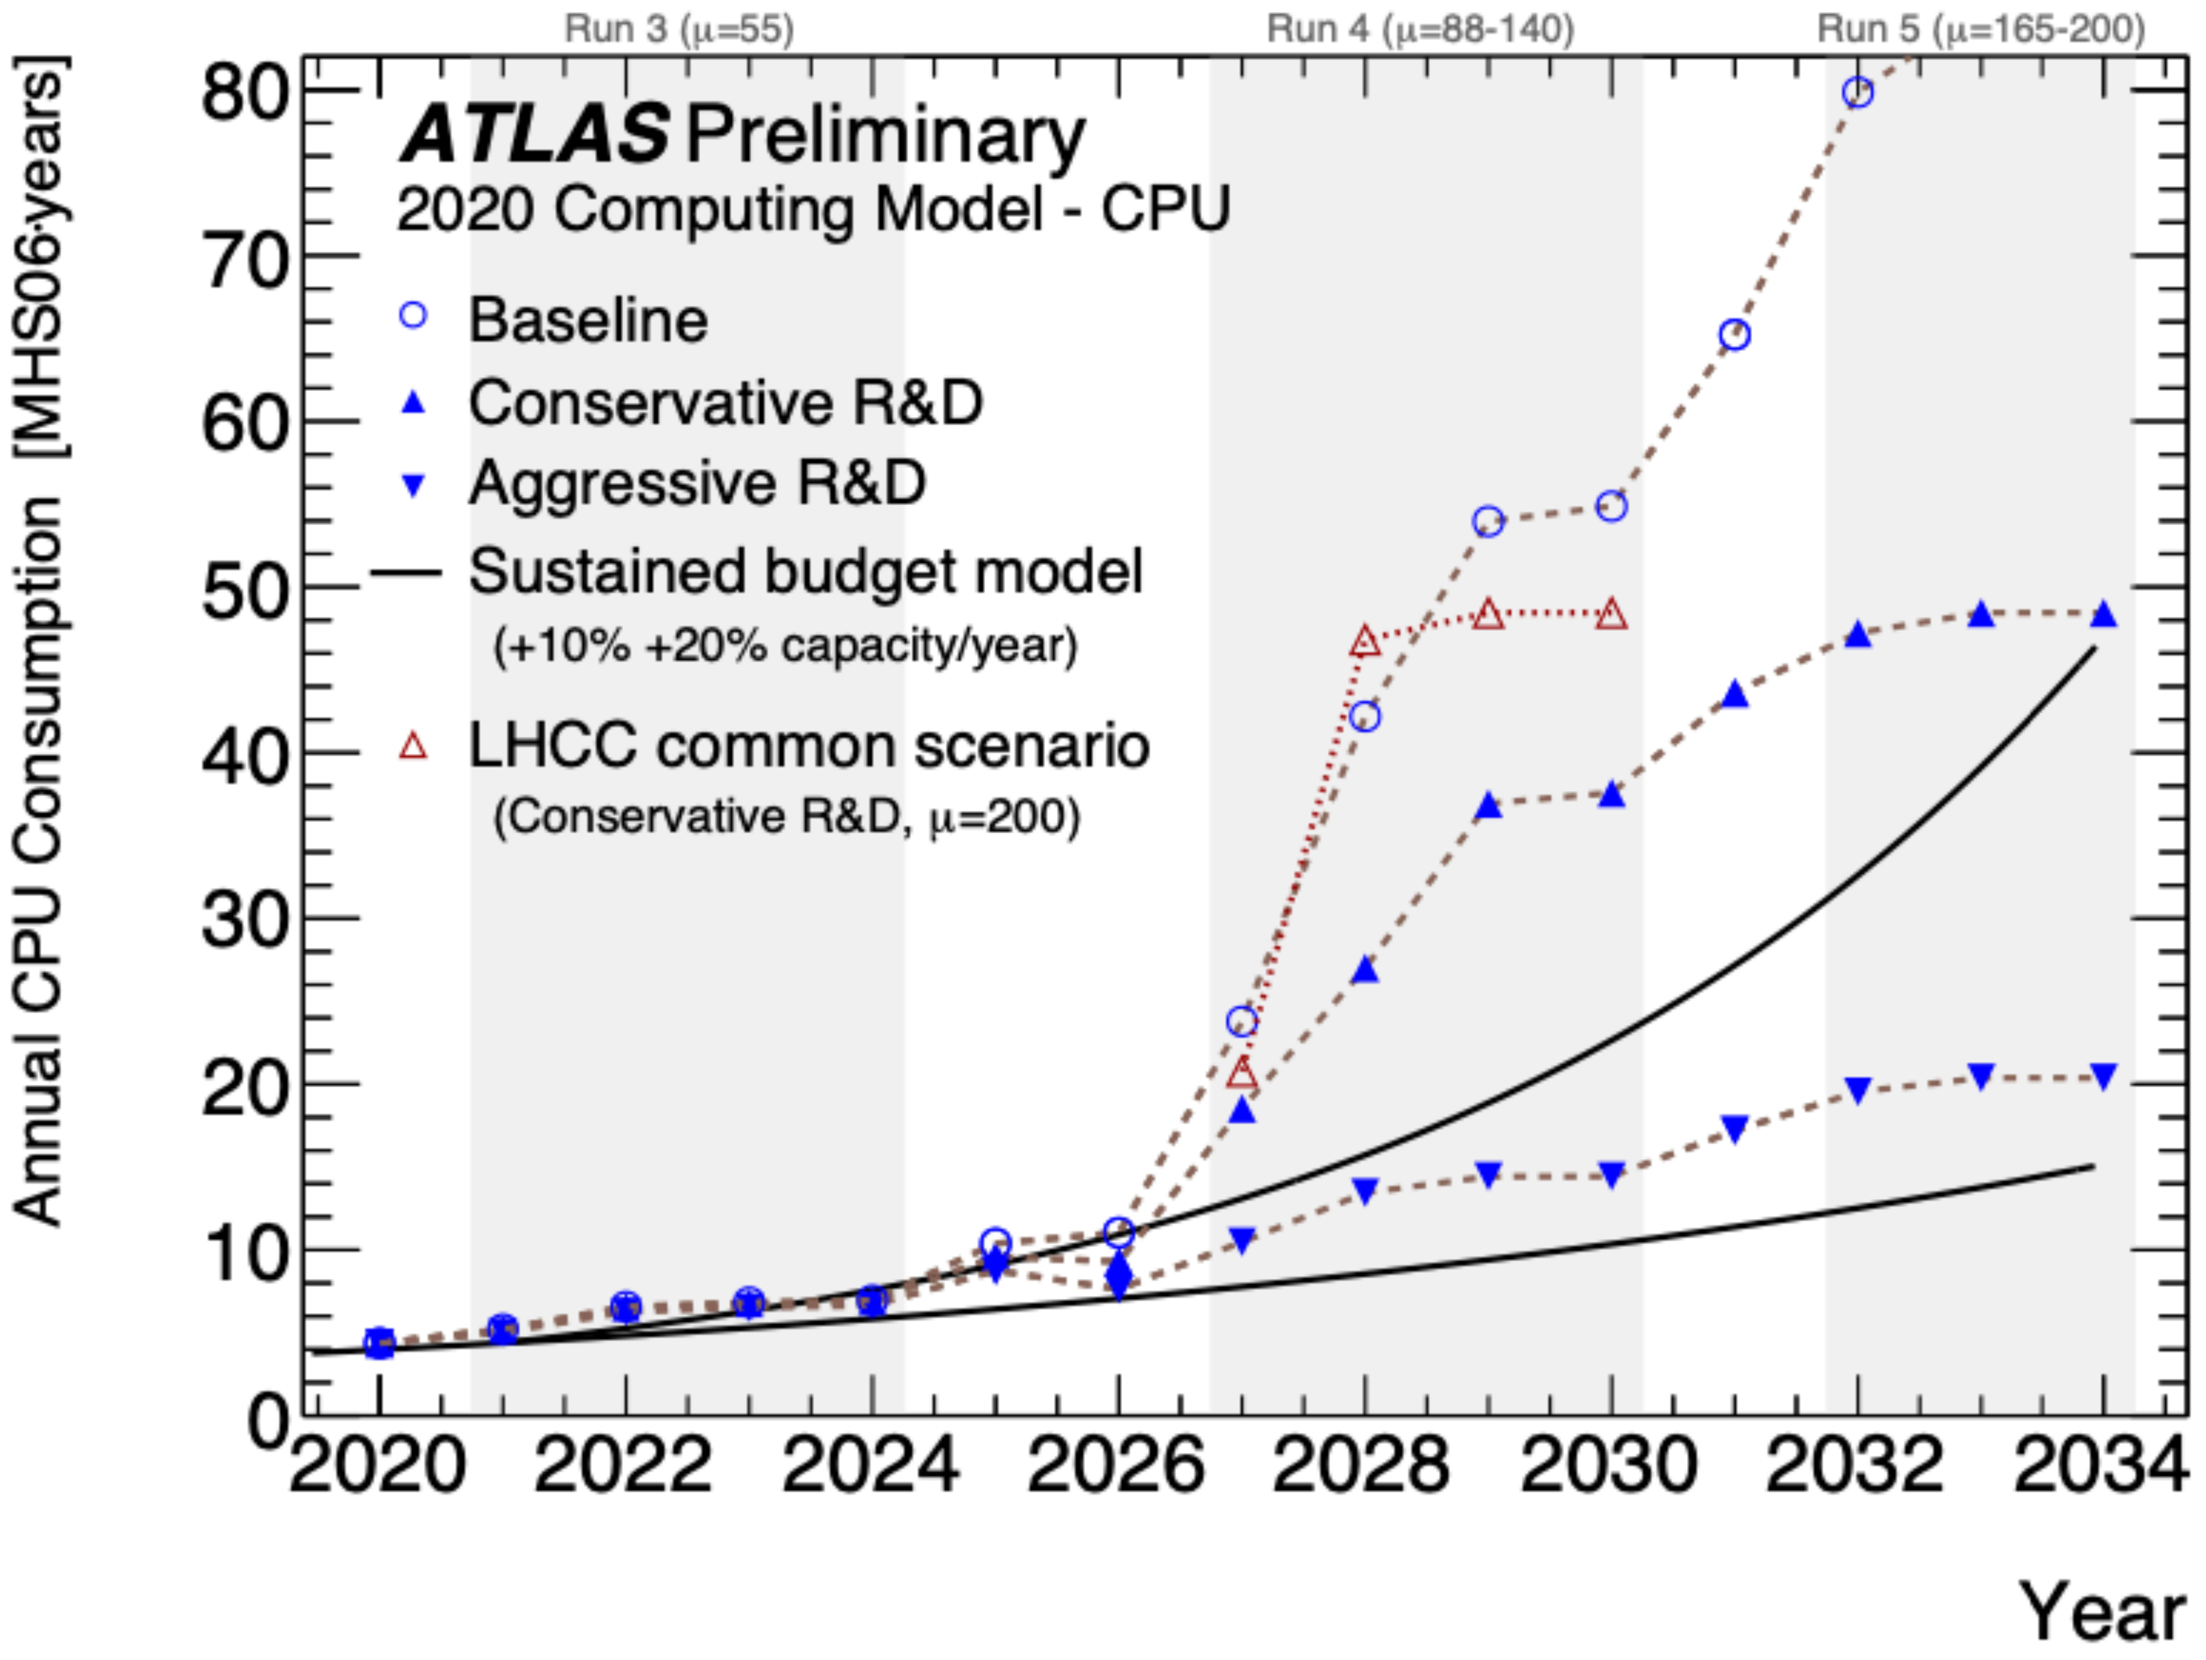

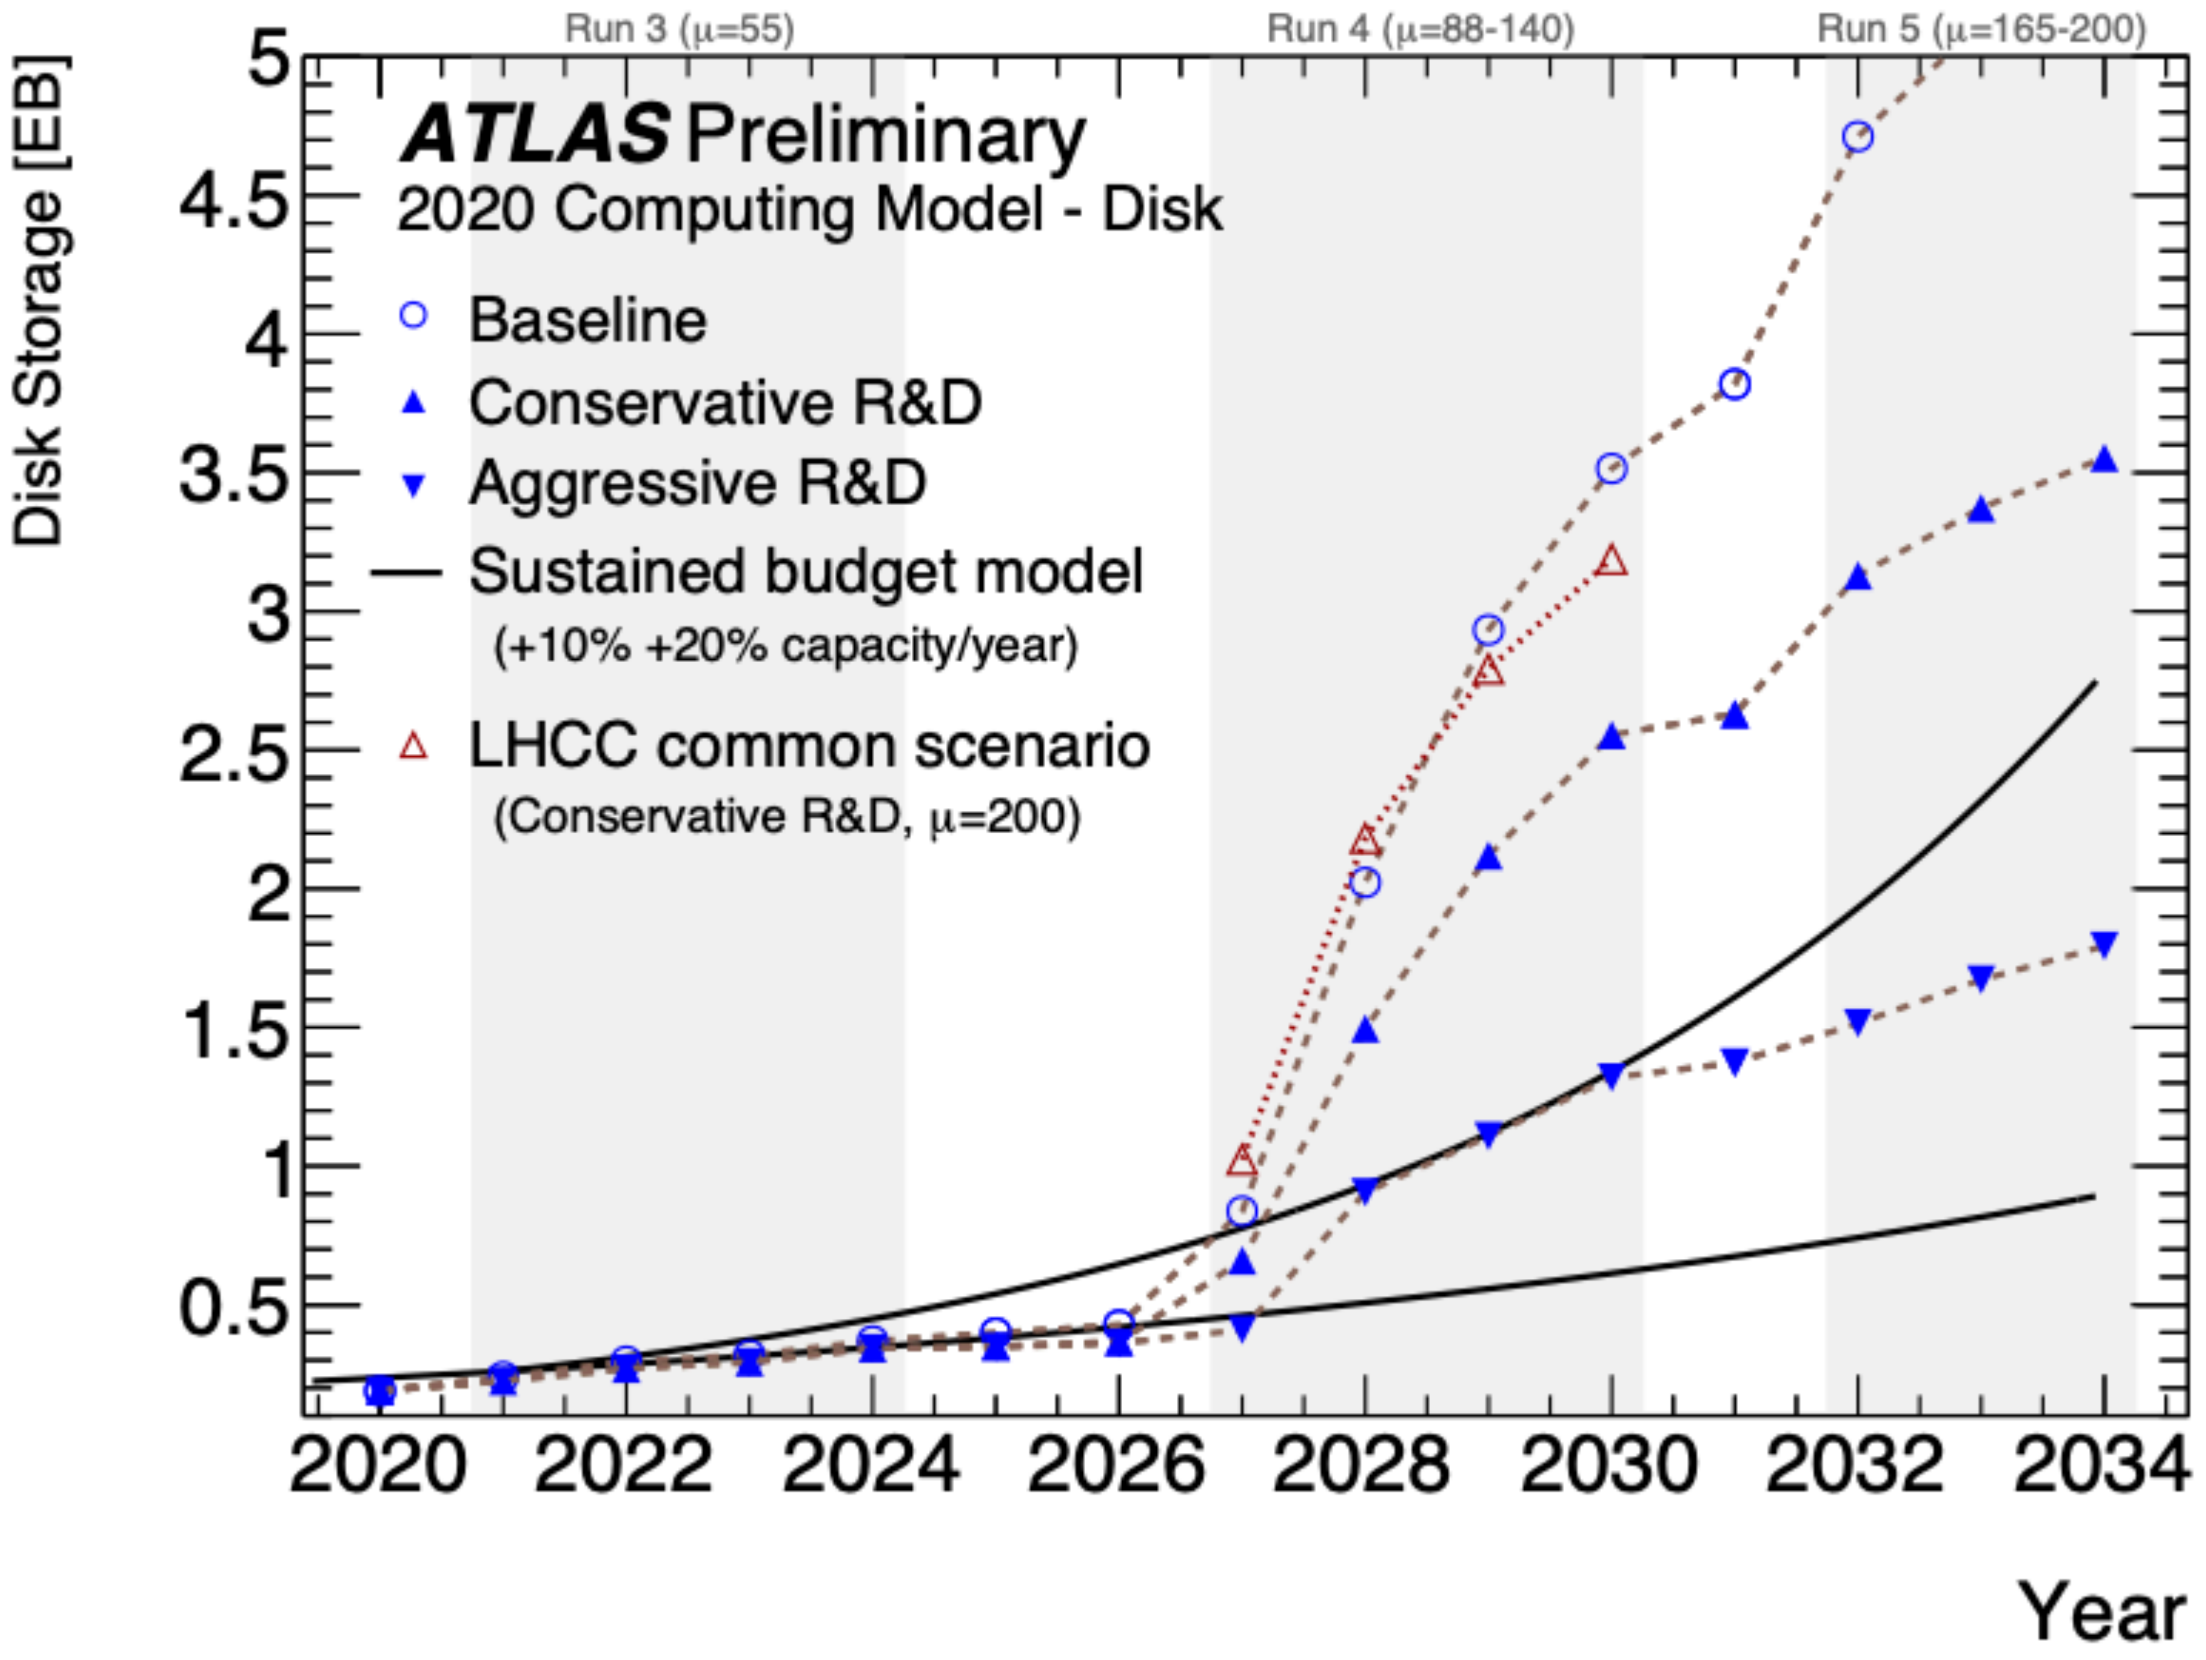

The outside world

… in the 2010s, many started facing the same problems as use and is finding their own solutions…


The New York Times

Artificial Intelligence › | A.I. Faces Quiz | How the A.I. Race Began | Key Figures in the Field | One Year of ChatGPT

Scientists See Promise in Deep-Learning Programs

Share full article

Deep Learning takes off

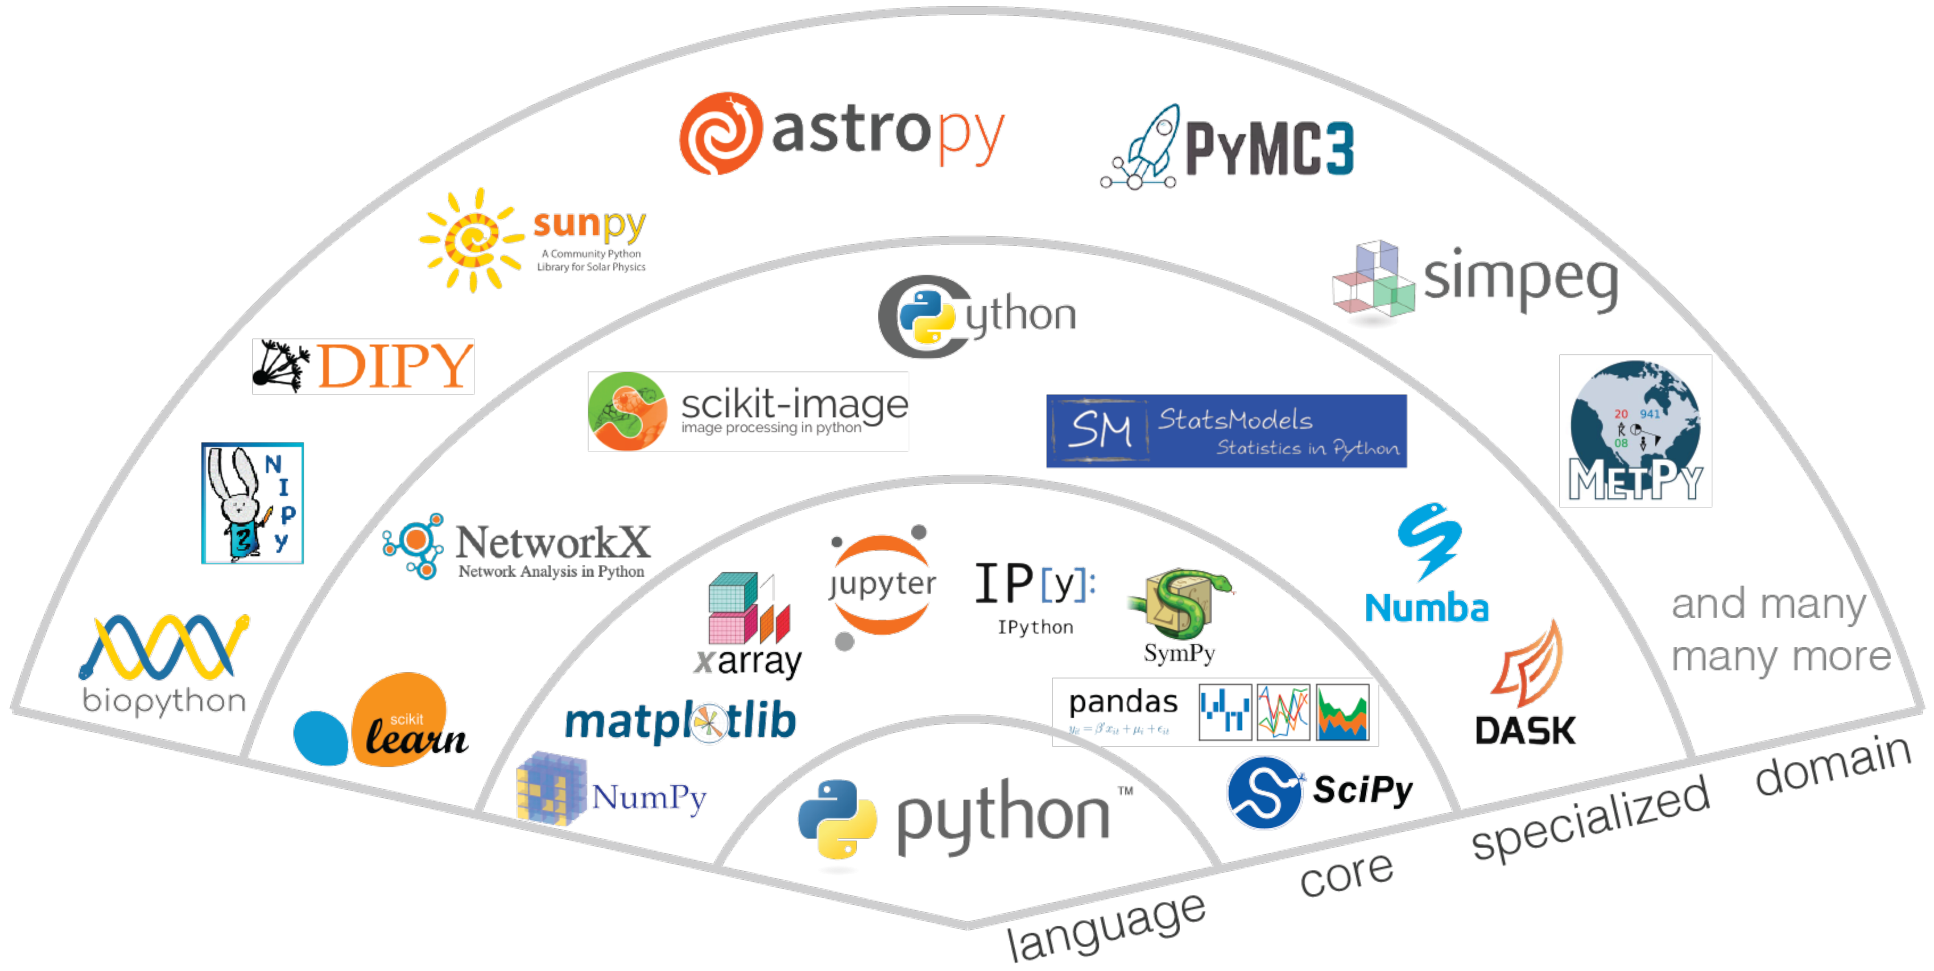

A new "data science" stack from the outside

kubernetes

openstack

MINIO

docker

A new distributed computing stack

The question

At HL-LHC users may deal with unprecedented data volumes for analysis. But analysis already now too slow, too cumbersome.

Made bearable by and for PhDs by paying for it with sweat & tears (babysitting jobs), cycles and storage (ntuples) → $$$

Can we imagine a different style of analysis than what we have now and if so, what would the infrastructure it look like that enables this for a large set of users?

“Analysis Facilities” Discussion

My personal interpretation: the more recent AF discussion is fueled by a sense that the time is ripe.

We have 20 years of LHC experience + new data analysis + infrastructure tools to realize new analysis patterns

→ HSF Analysis Facilities Forum

HEP Software Foundation

HSF Analysis Facilities Forum

Mandate: Forum to discuss recent developments on R&D related to new ideas on analysis infrastructure

Produce report summarizing trends & observations of current R&D

Conveners:

Alessandra Forti
ATLAS

Nicole Skidmore
LHCb

Diego Ciangottini
CMS

LH
ATLAS

User Perspective Requirements

Ability to perform fast research iterations on large datasets interactively

Ability to convert interactive to batch-schedulable workloads

Ability to interact with the WLCG and scale outside of the facility on occasion

Ability to efficiently train machine learning models for HEP

Ability to reproducibly instantiate desired software stack

Ability to collaborate in a multi-organisational team on a single resource

Ability to move analyses to new facilities

Ability to efficiently access collaboration data & make intermediate data products available

Ability to express interdependent distributed computations at small and large scales

Beyond ssh - Jupyter & etc as Entry Point

Jupyter should be thought of as a “richer shell” i.e. similar to sessions at login nodes. **I should be able to address / work with scaleout backend from within such a notebook (incl. batch system)**

Coexist peacefully with e.g. terminal based access

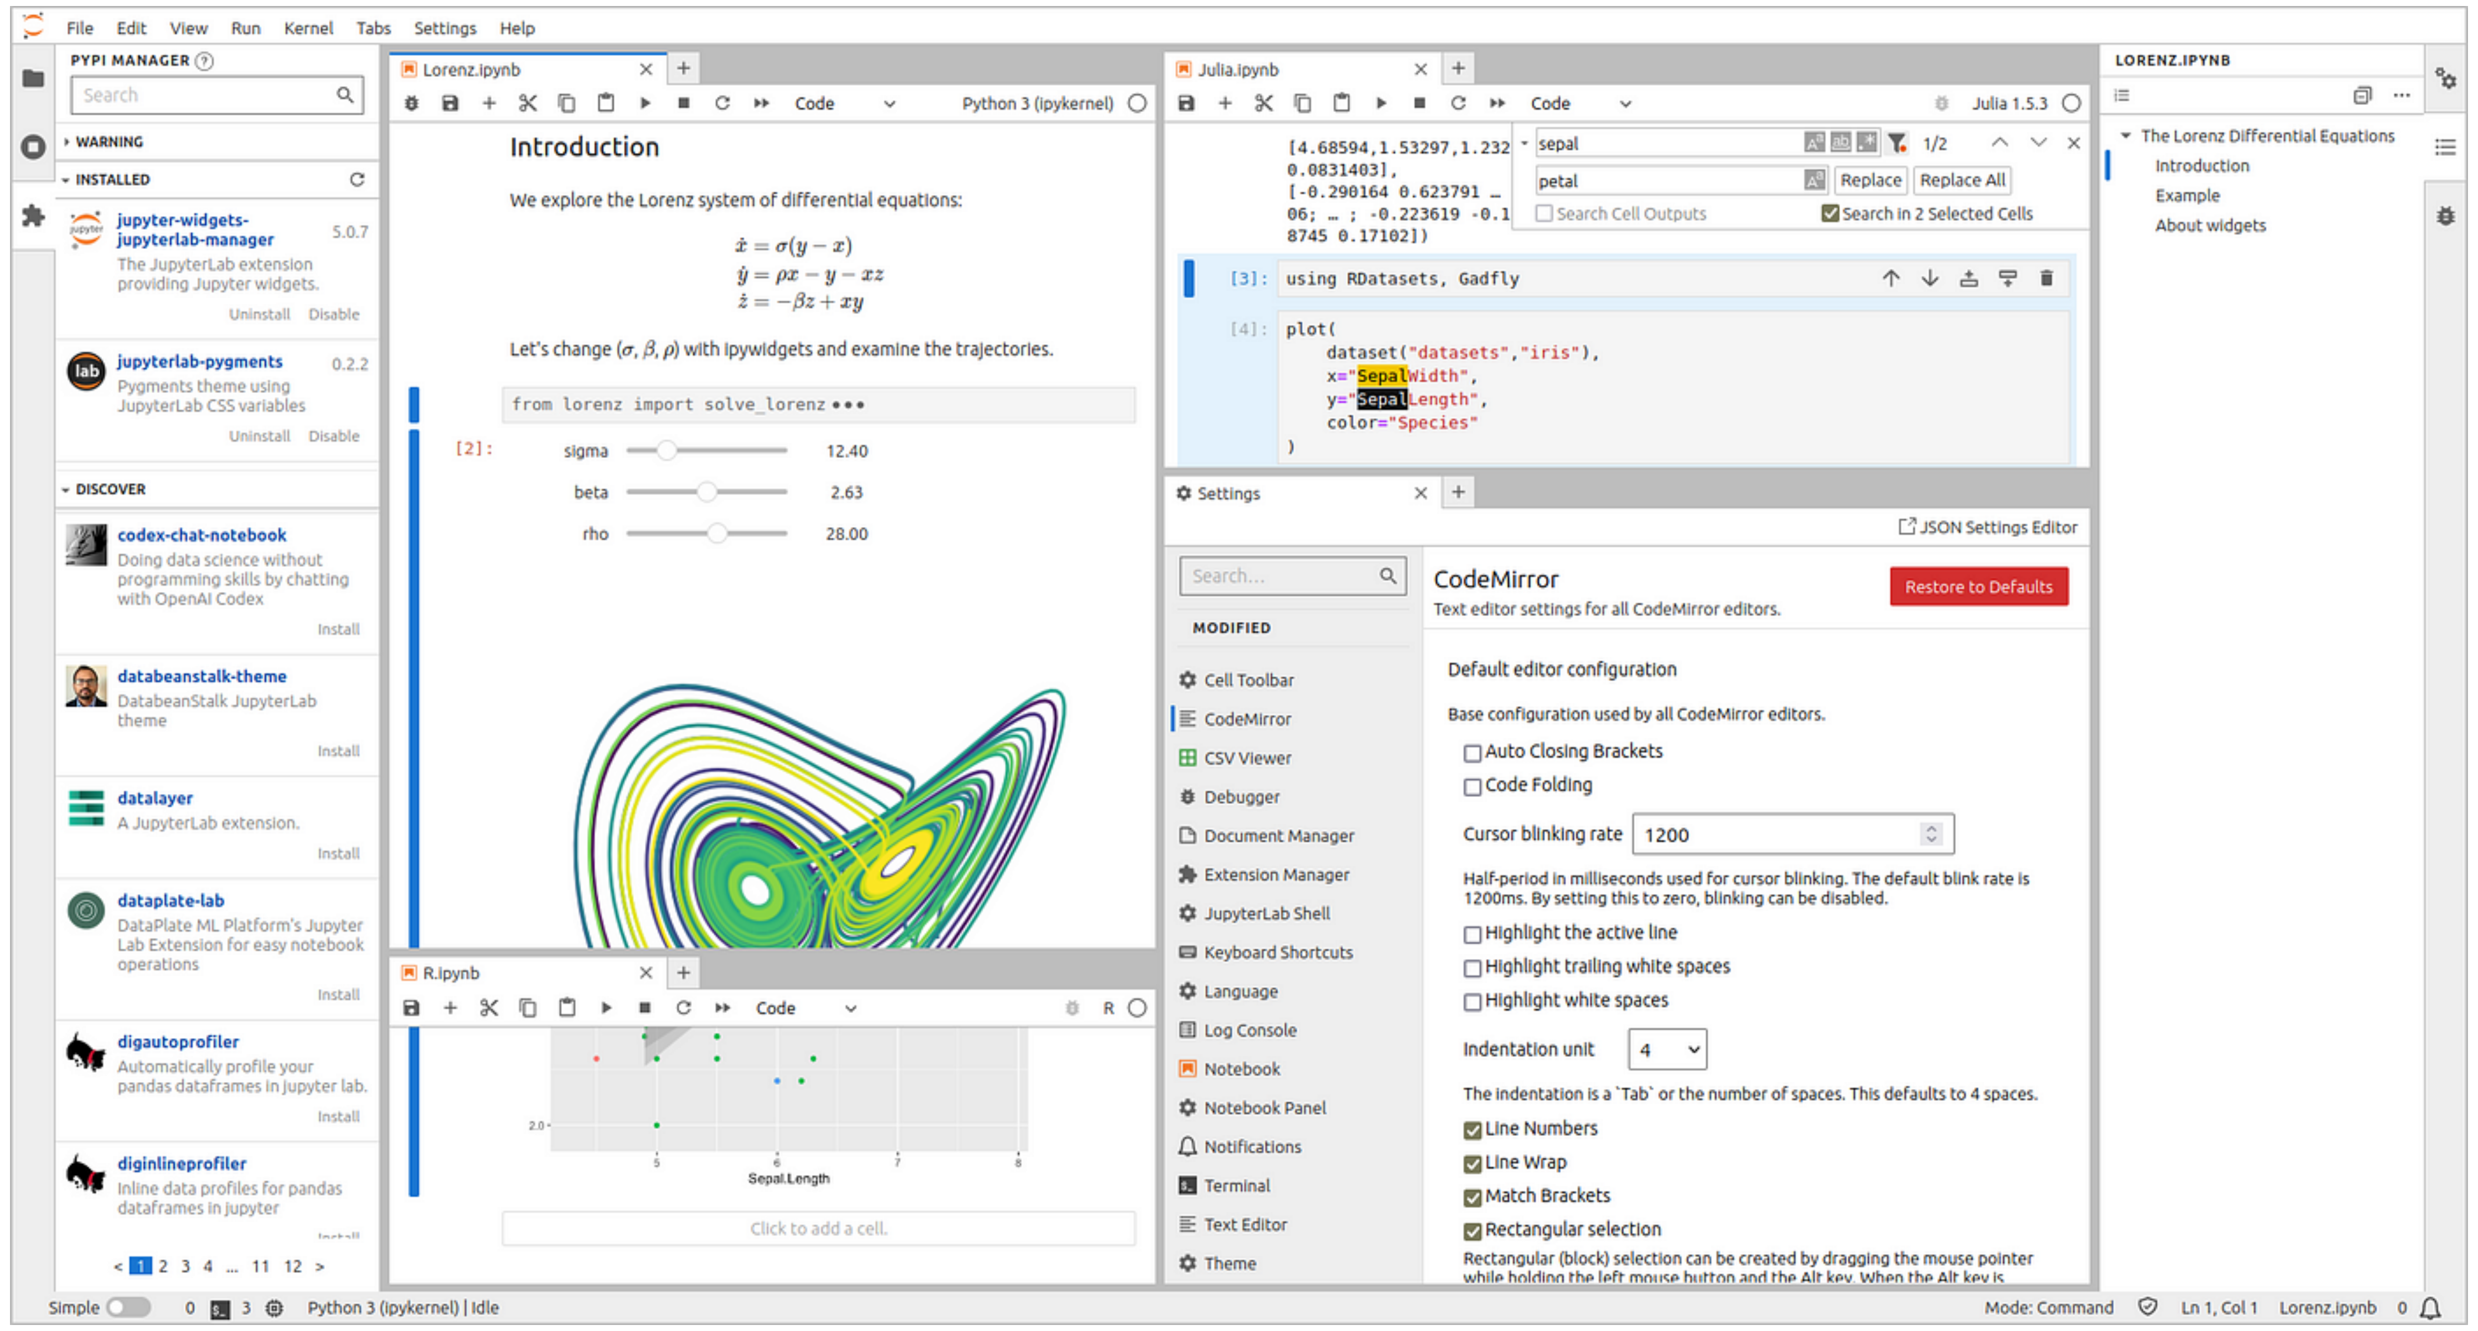

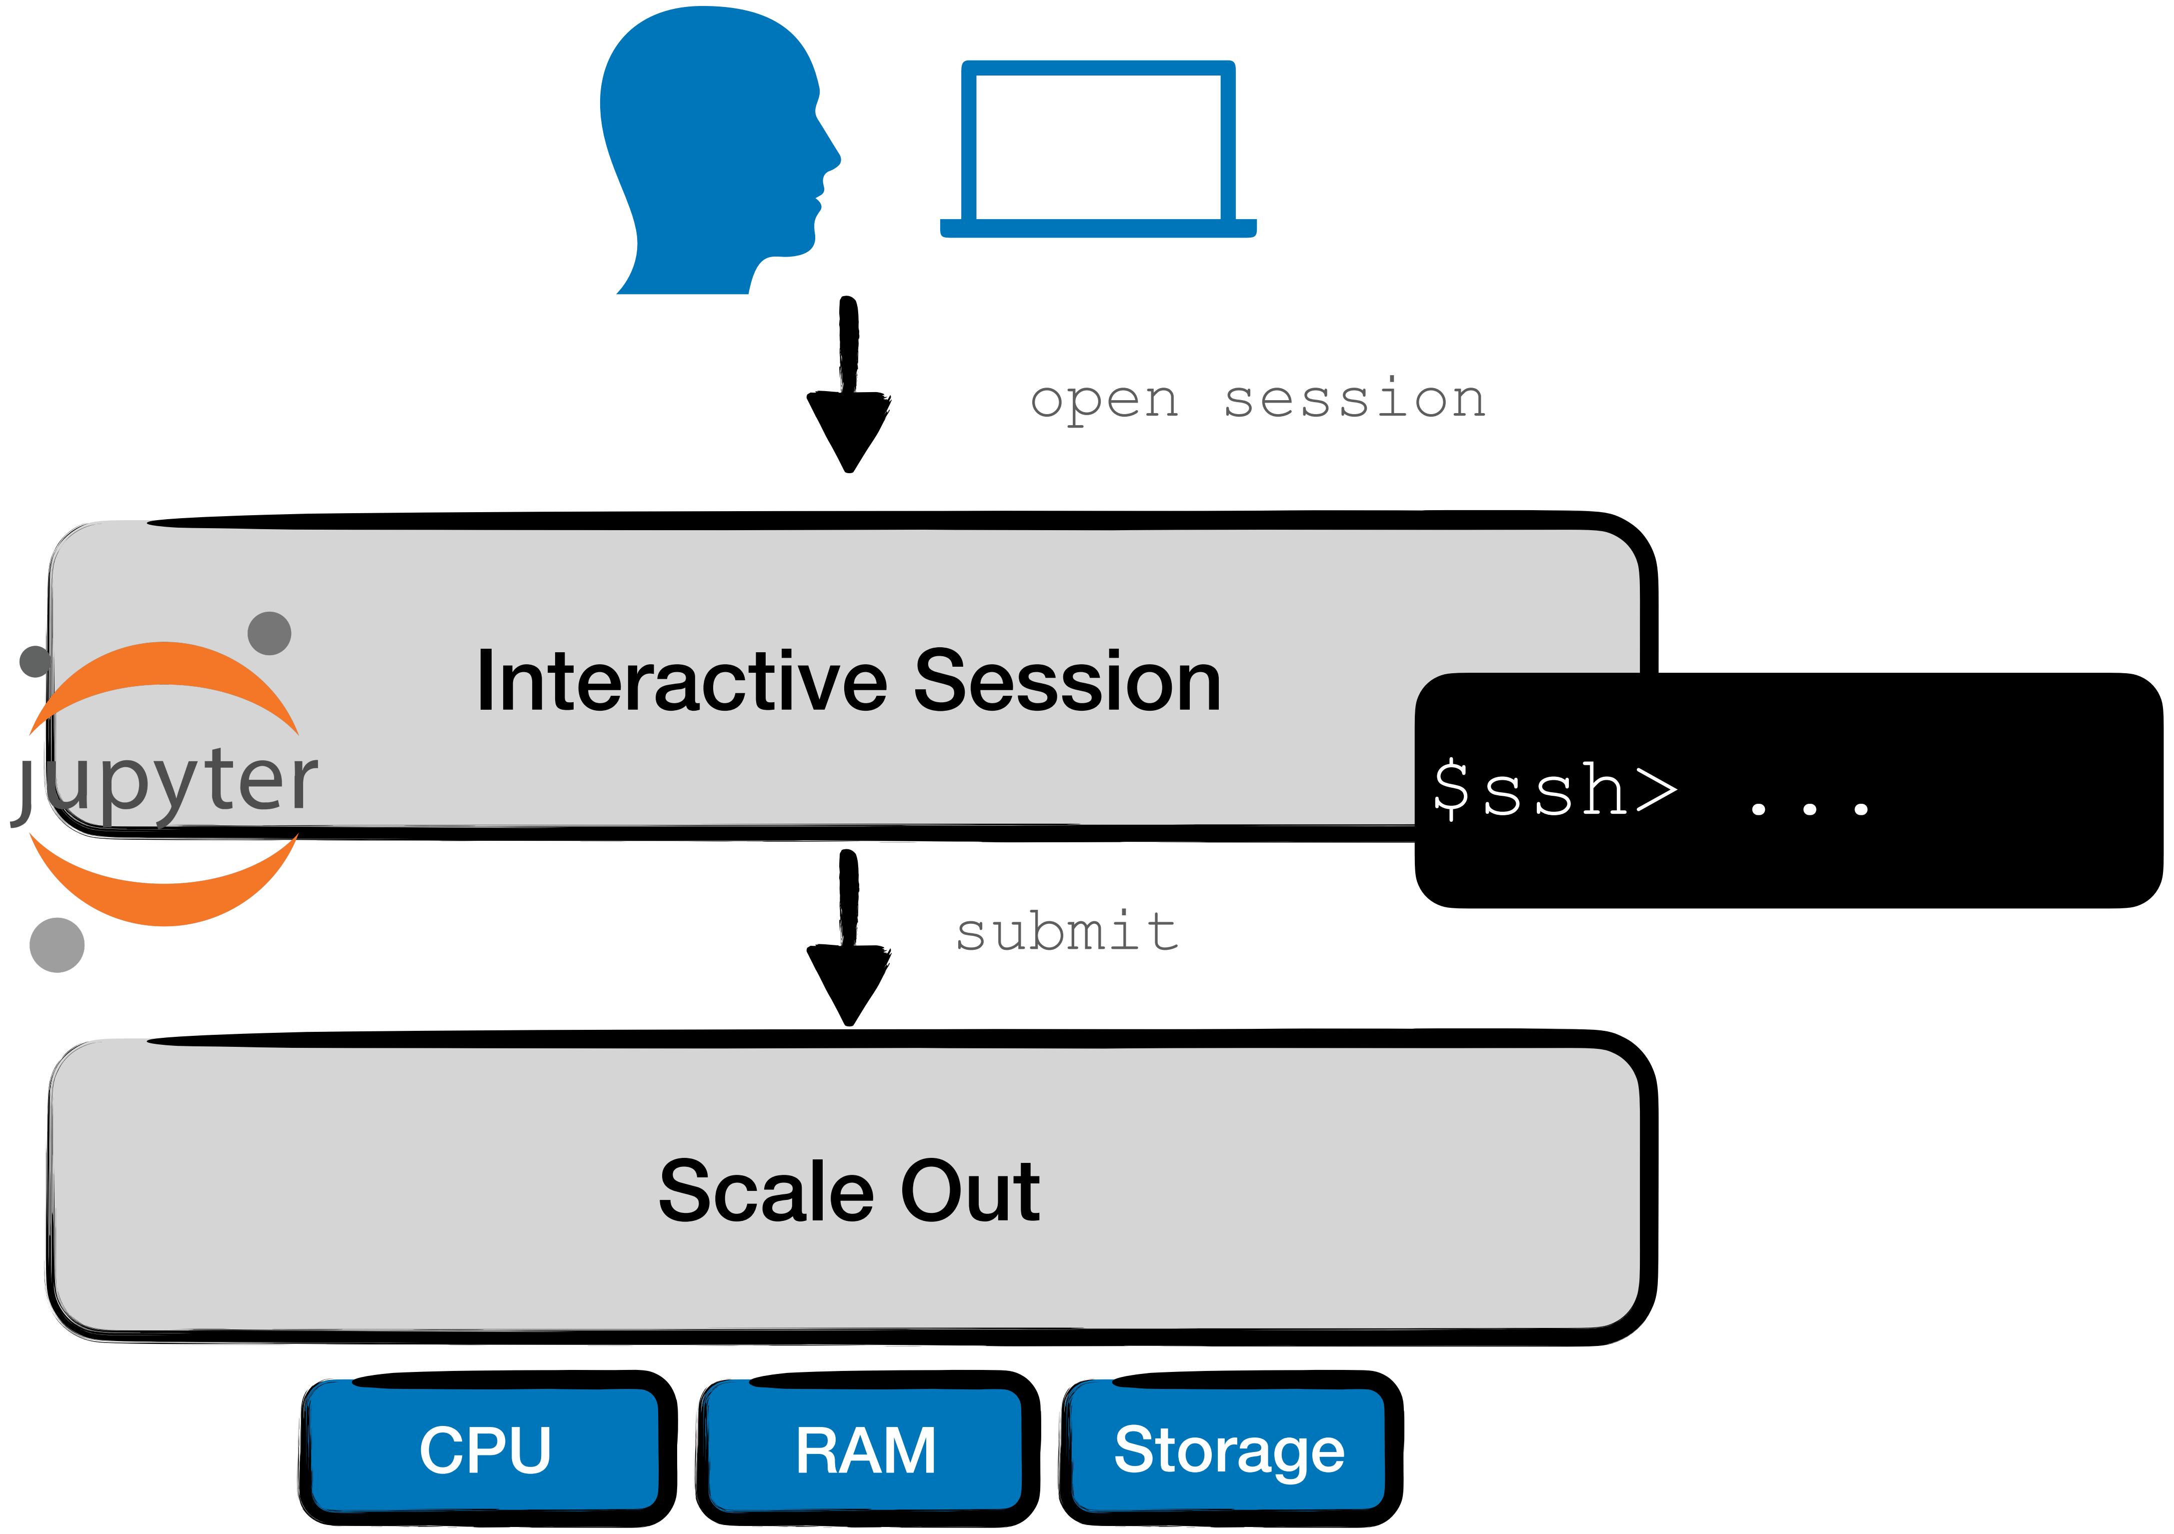

Interactive Distributed Scale-Out

From the user's perspective, the hope for a method to run interactive scale-out with fast turnaround times drives AF R&D

This is true in both "ecosystems": **ROOT** and **Python**

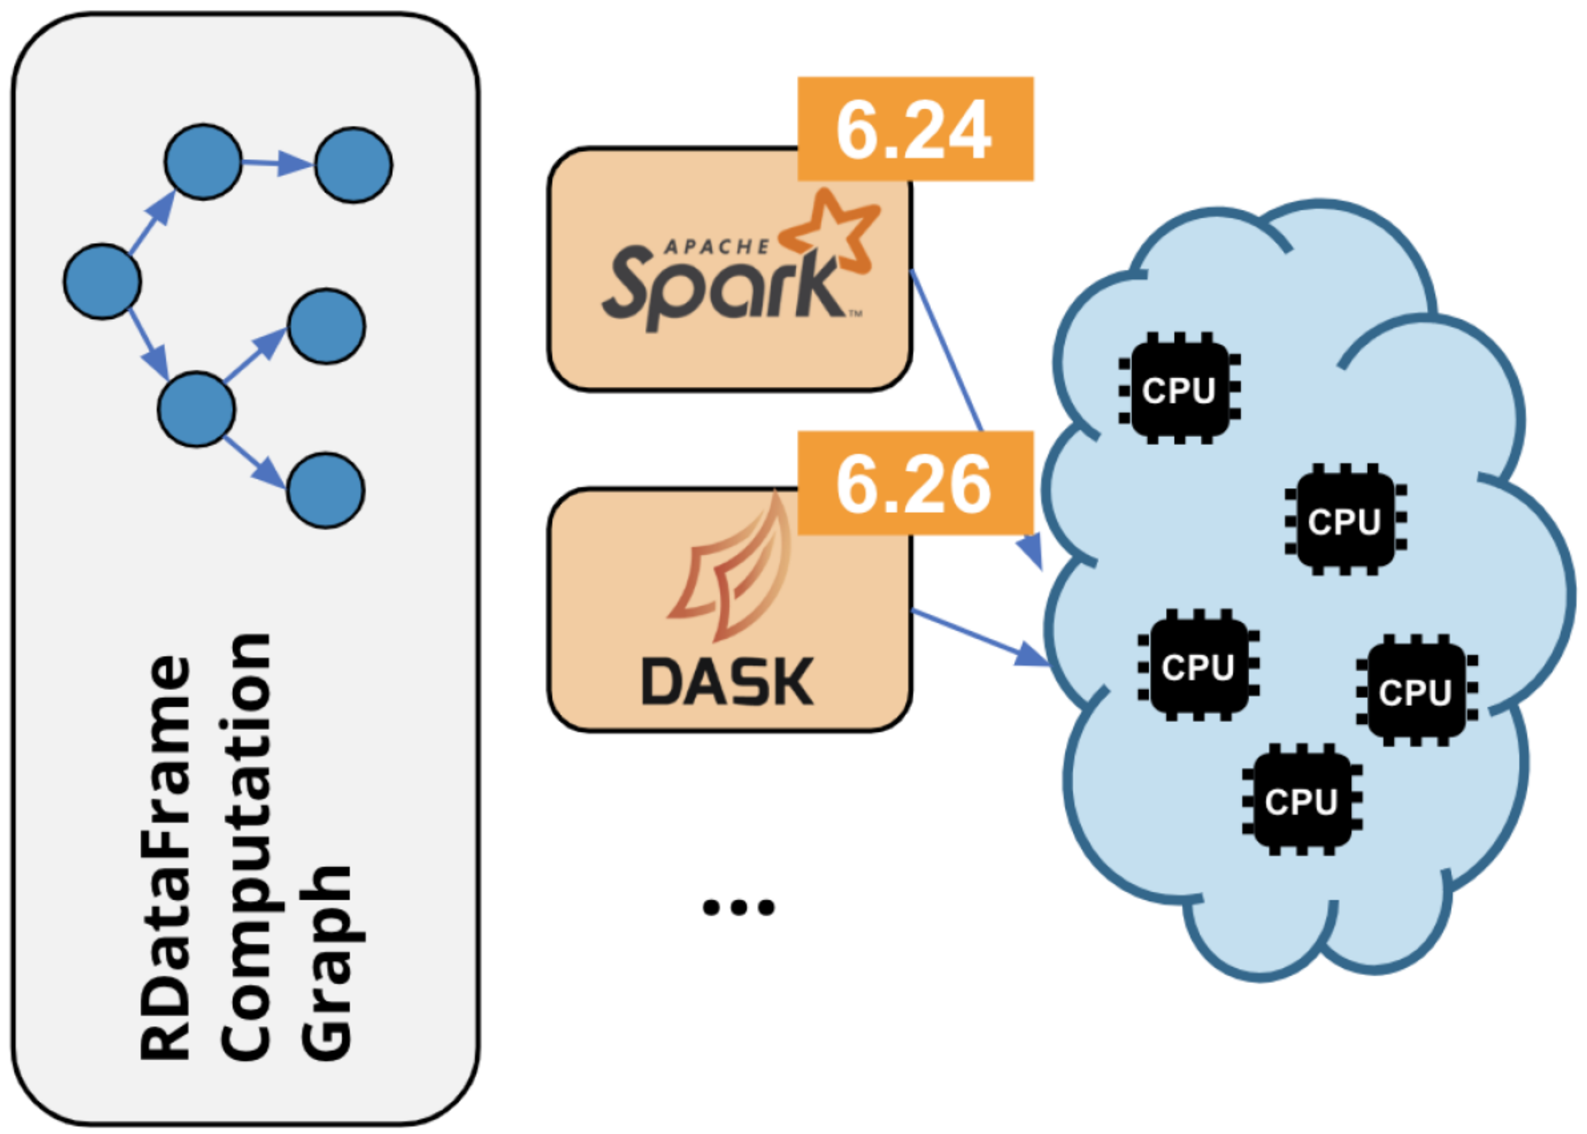

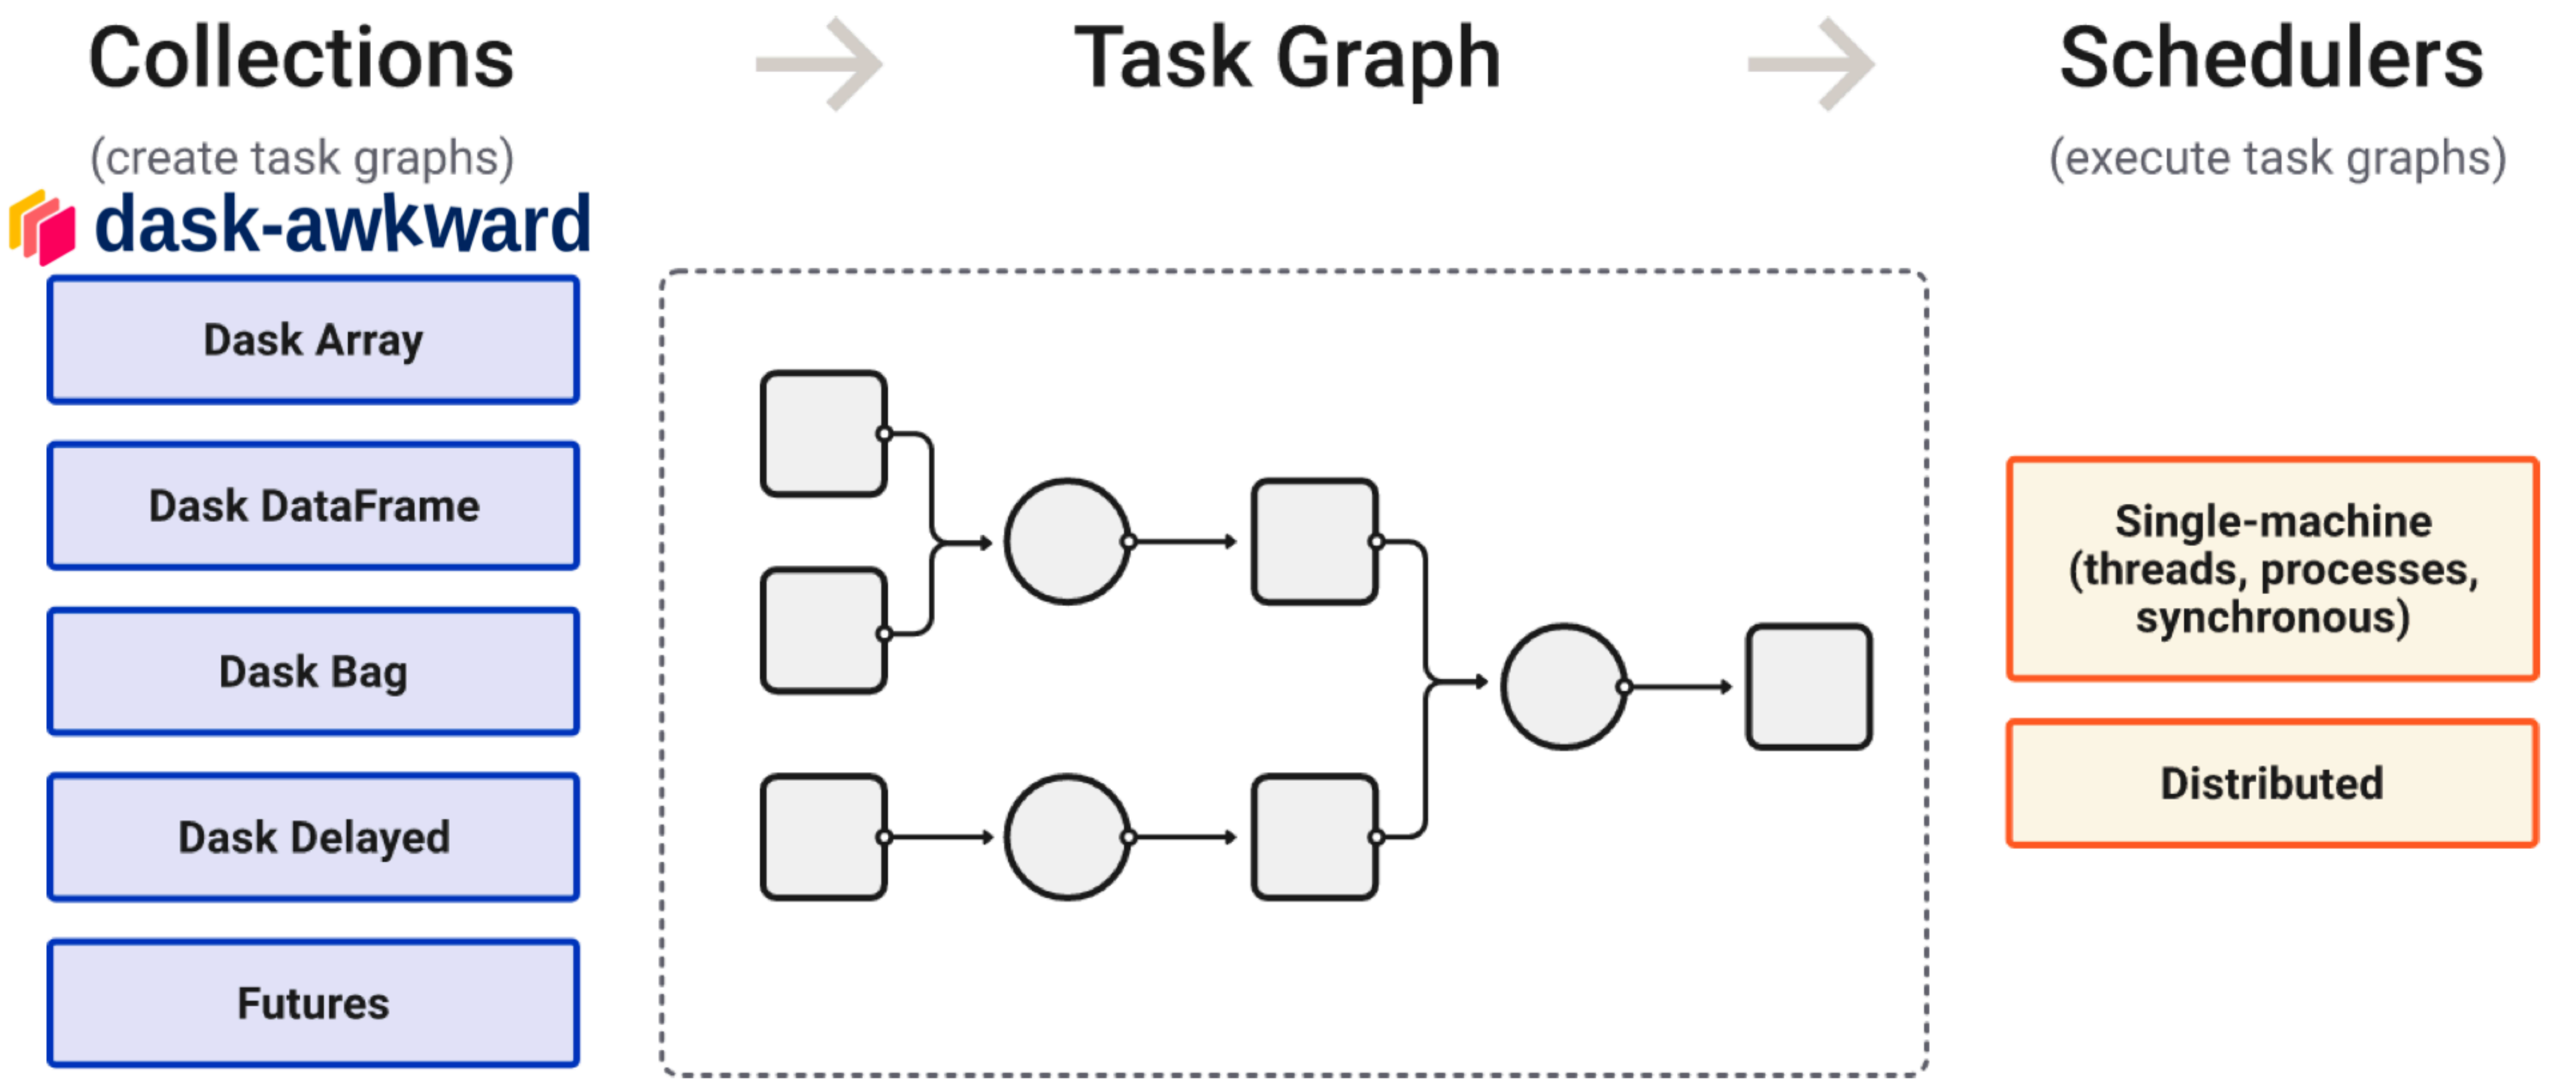

ROOT

scikit-hep based (coffea, awkward, uproot, …)

Convergence (?) of technologies

Both ecosystems seem(?) to have picked **Dask** as the primary scale-out system. **Extending AFs** such that multiple users can reliably launch, operate & scale Dask clusters is a common trait.

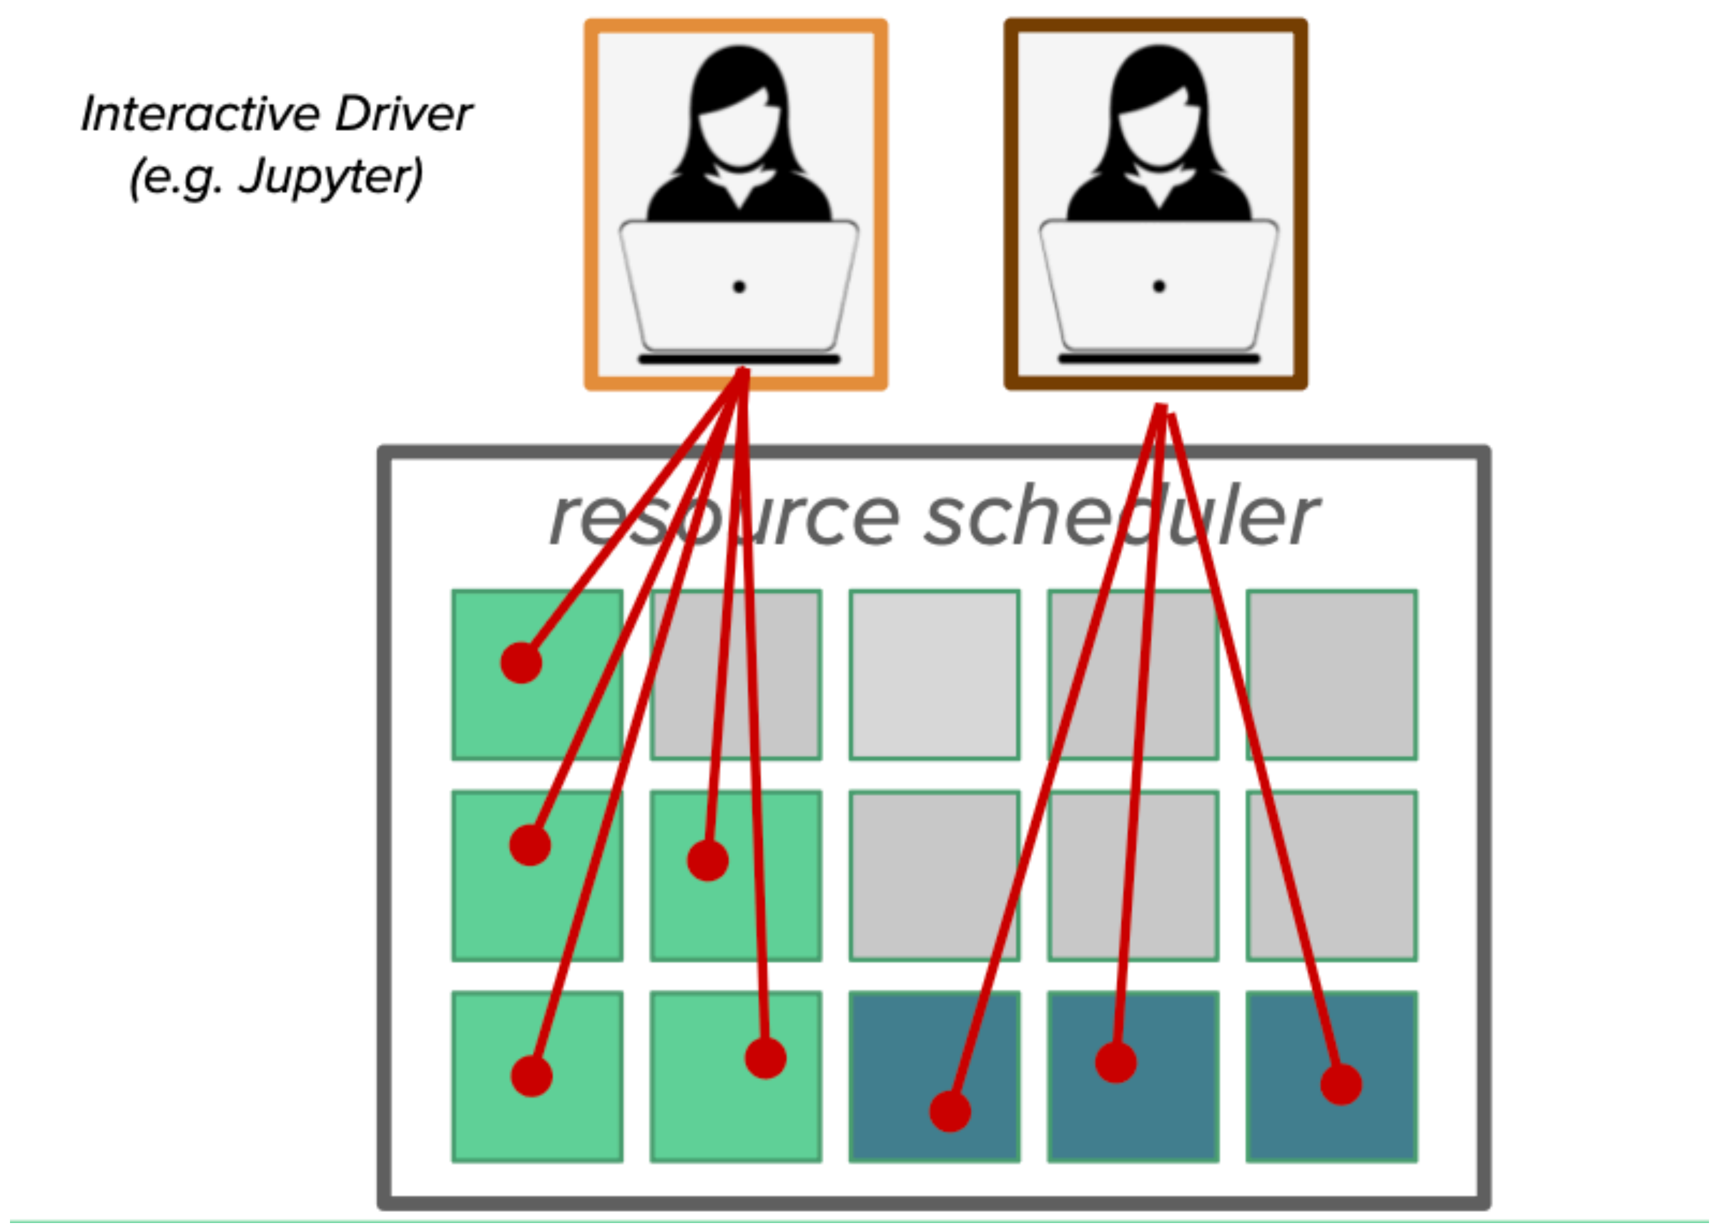

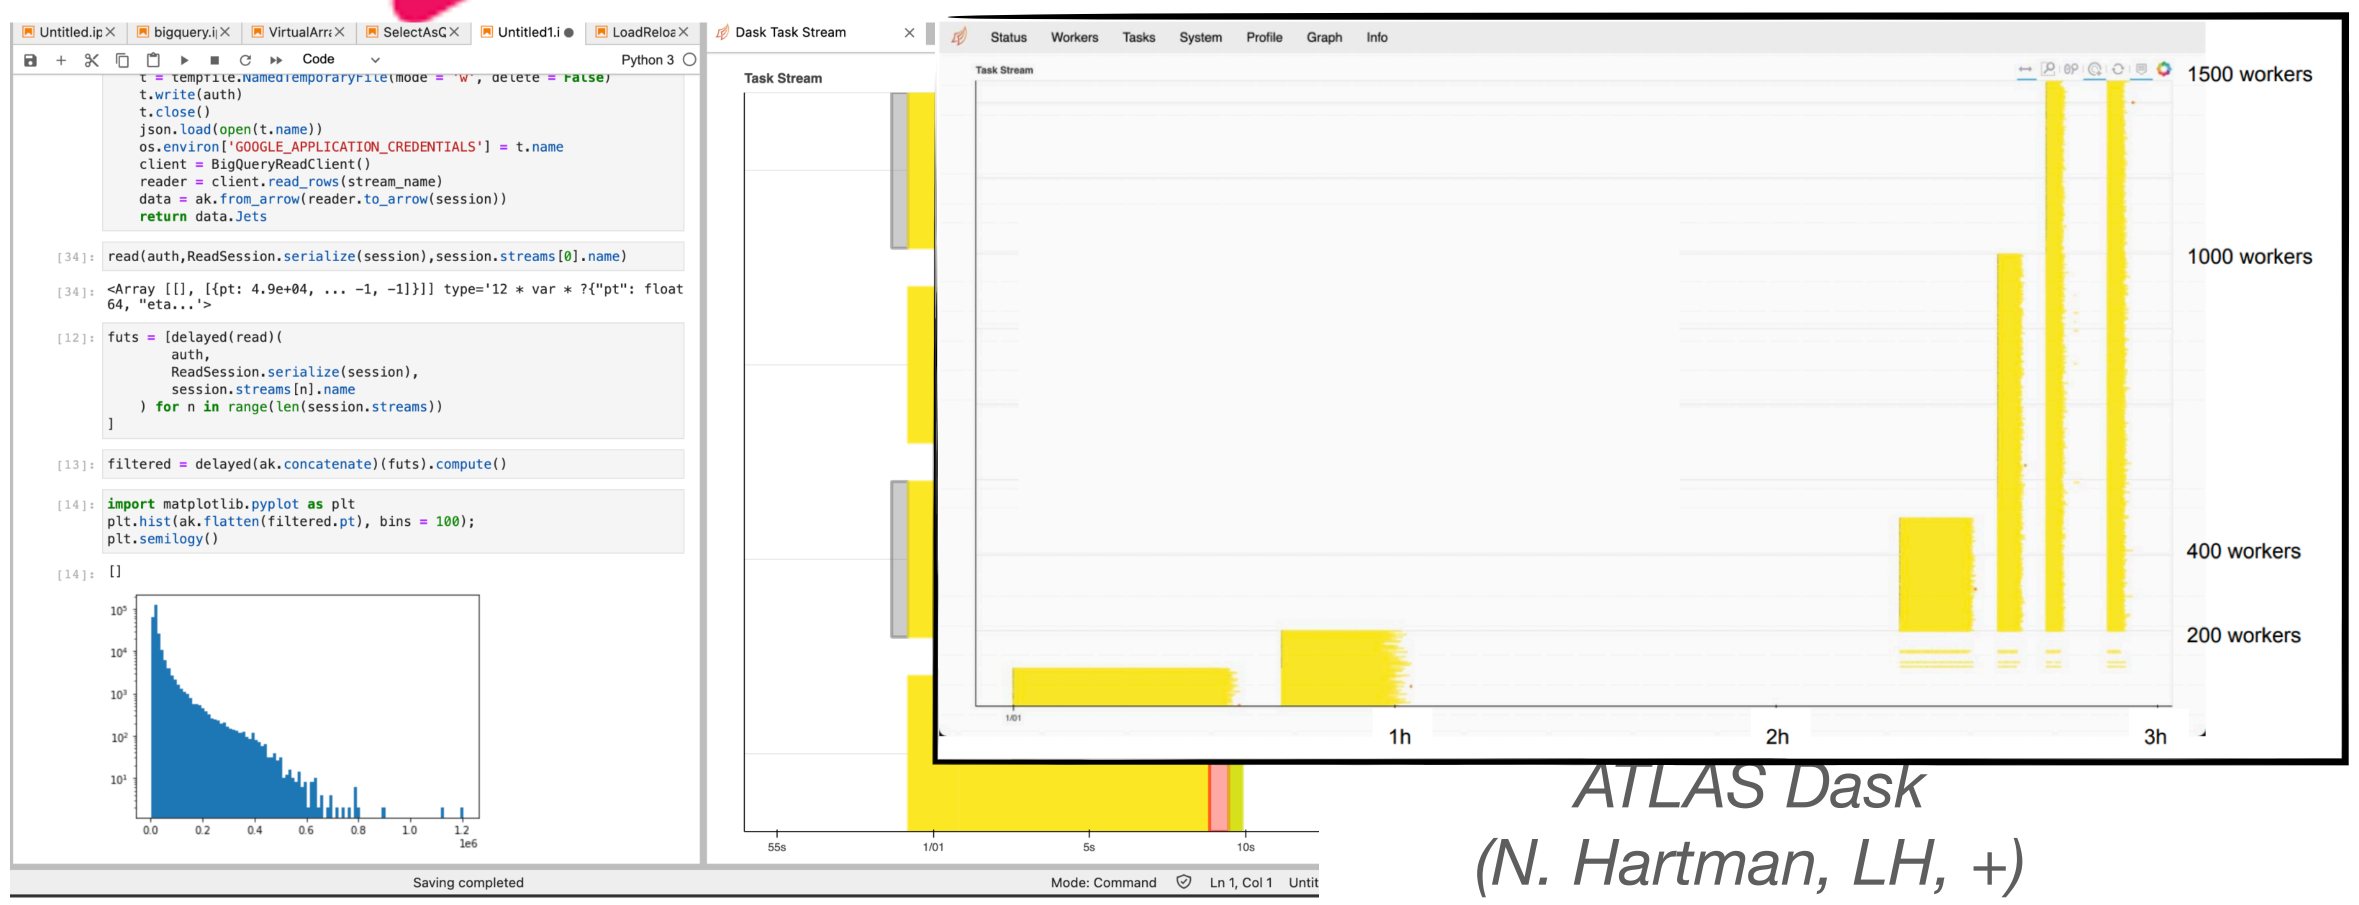

ATLAS Dask
(N. Hartman, LH, +)

HL-LHC Readiness

The AF development is closely linked also to efforts to benchmark HL-LHC. Focus on User Experience (columnar analysis, etc..)

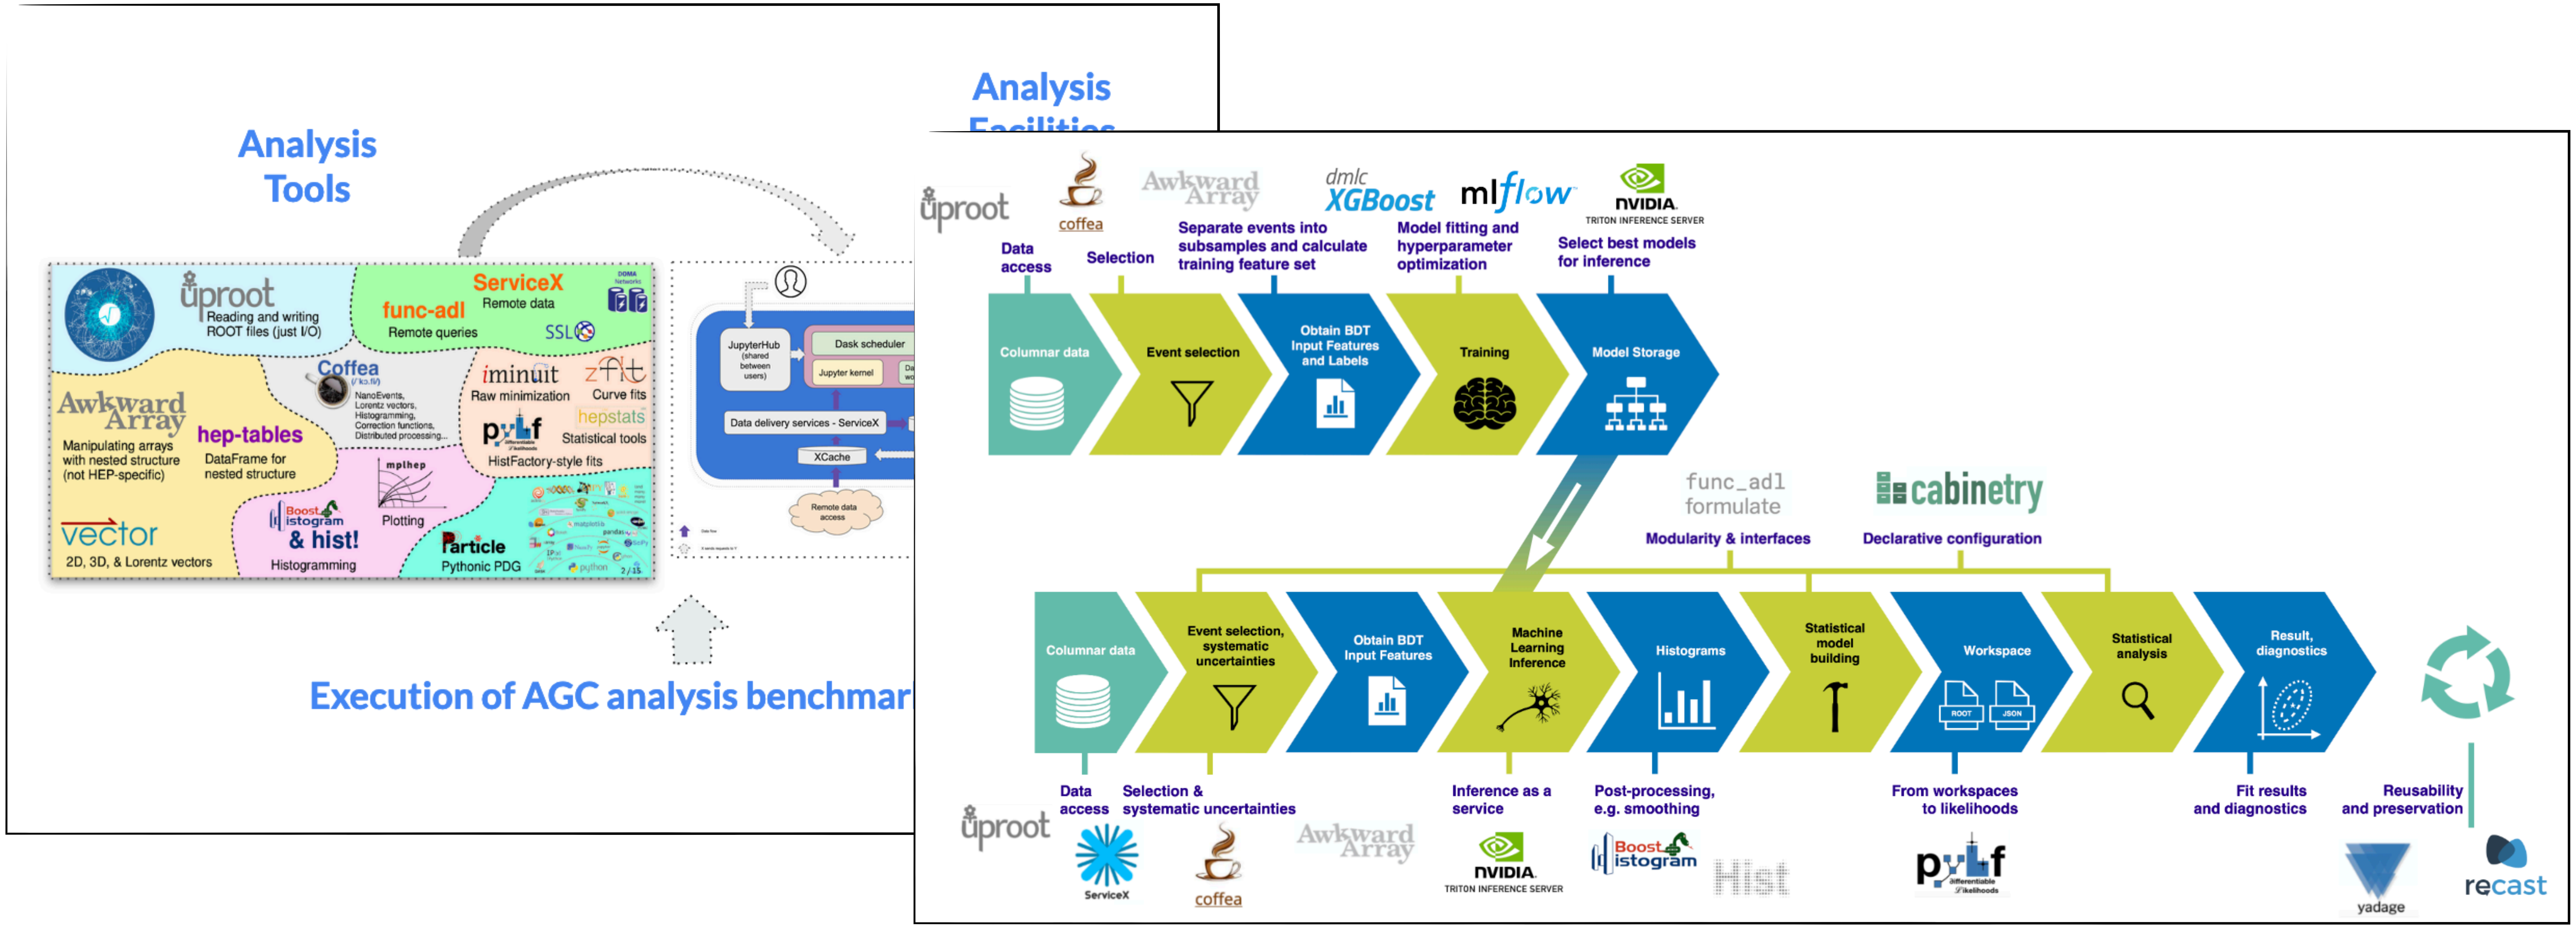

Headless Scale-out

Interactive is nice, but just because I want to use dask, doesn't mean I want it to be interactive. Need solutions that allow to submit "distributed dataframe" workflows non-interactively

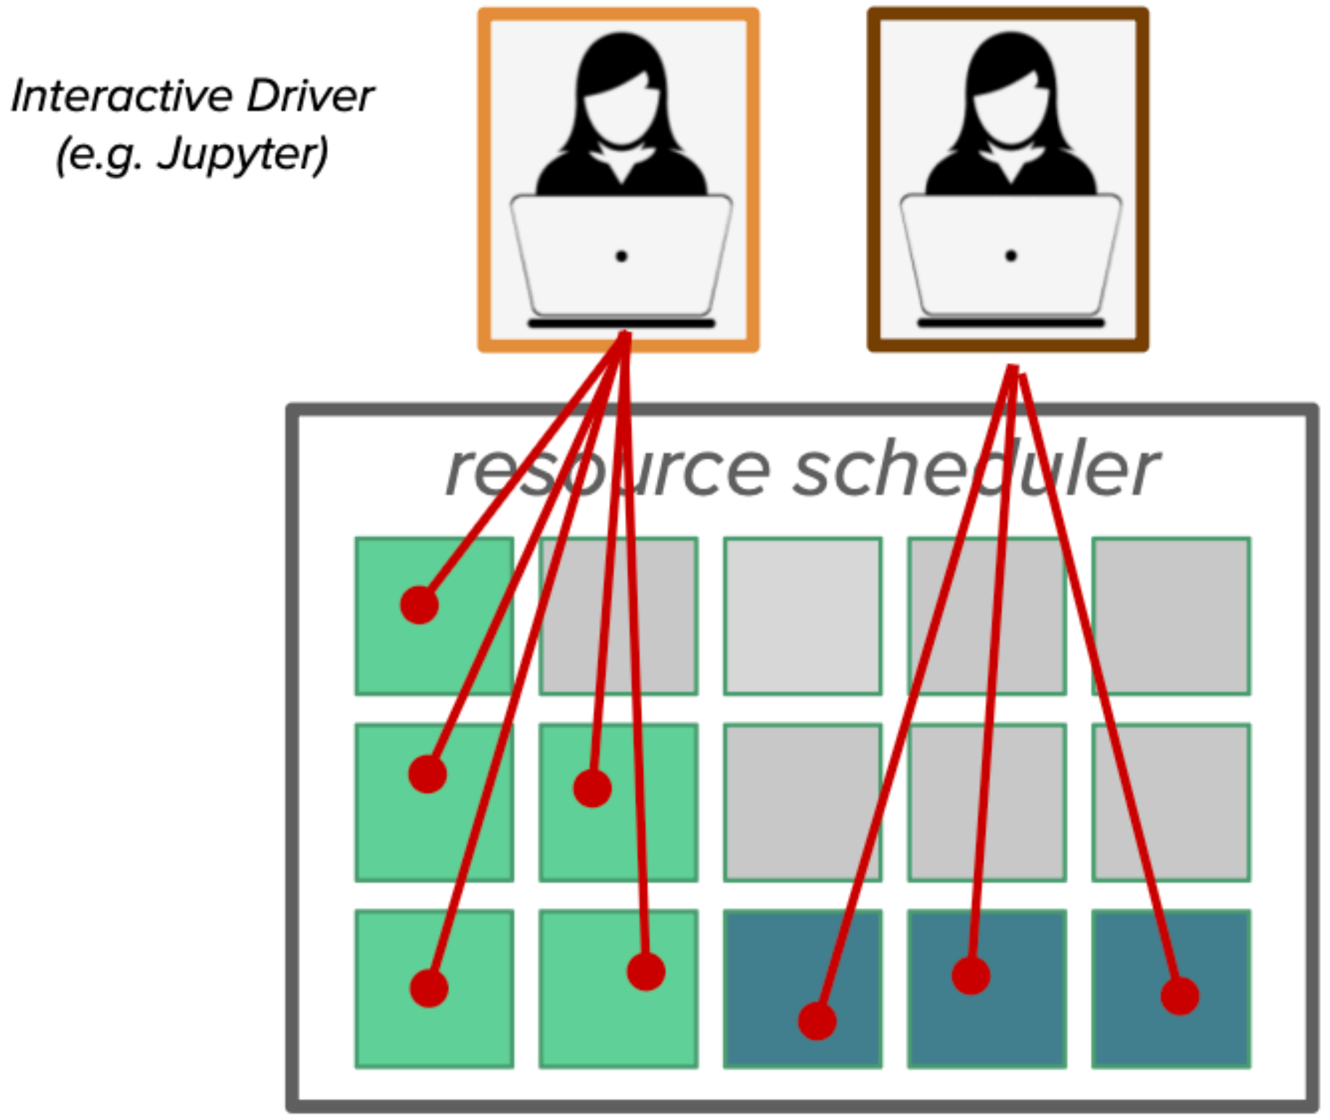

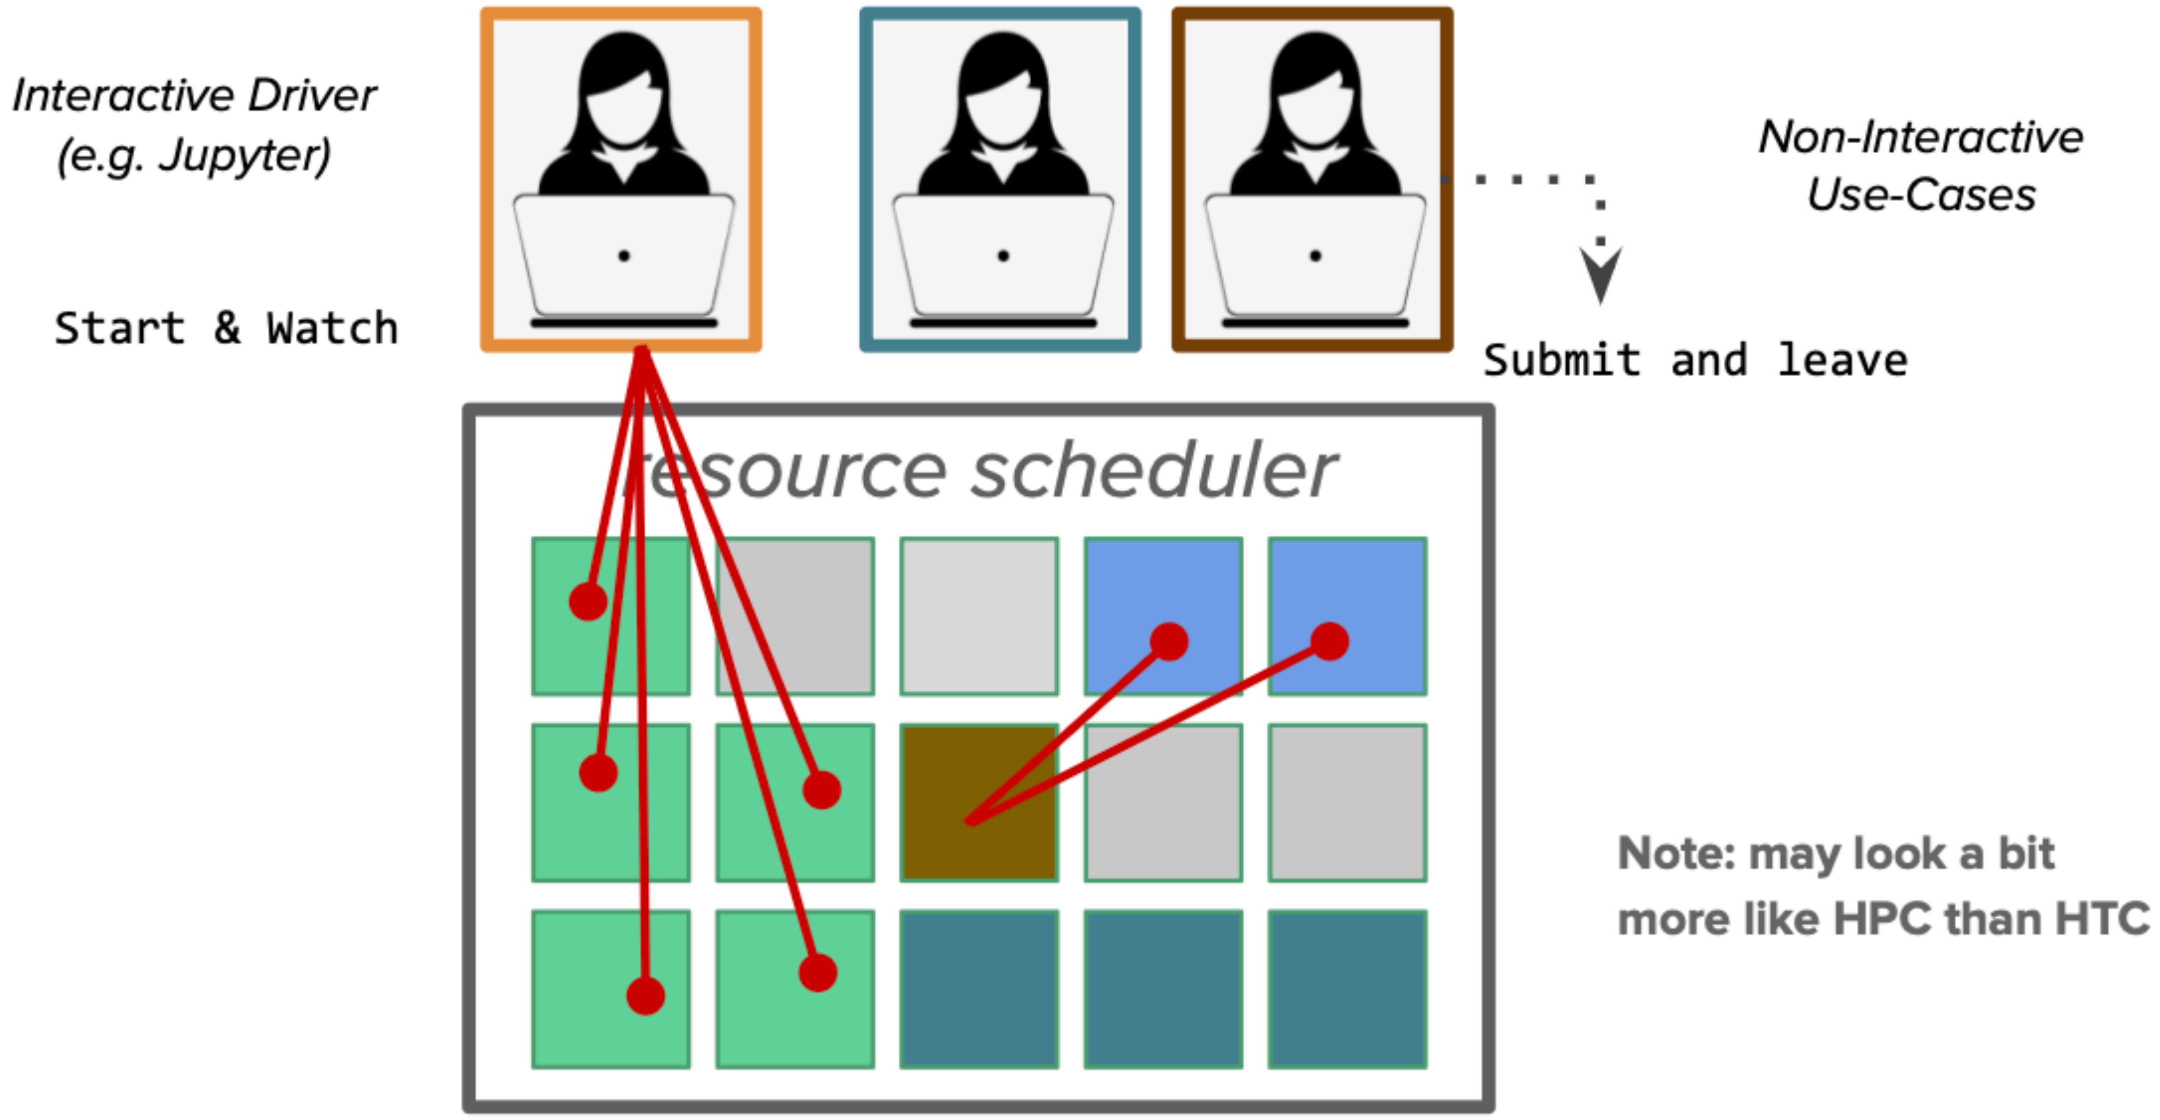

Machine Learning & GPUs

ML will only get a more prominent role in HEP analysis and requires / leads to very different workflows

- Data Exploration, Interactive R&D and small-scale training
- Large-scale non-interactive training and HP optimization
- ML Inference within an analysis pipeline

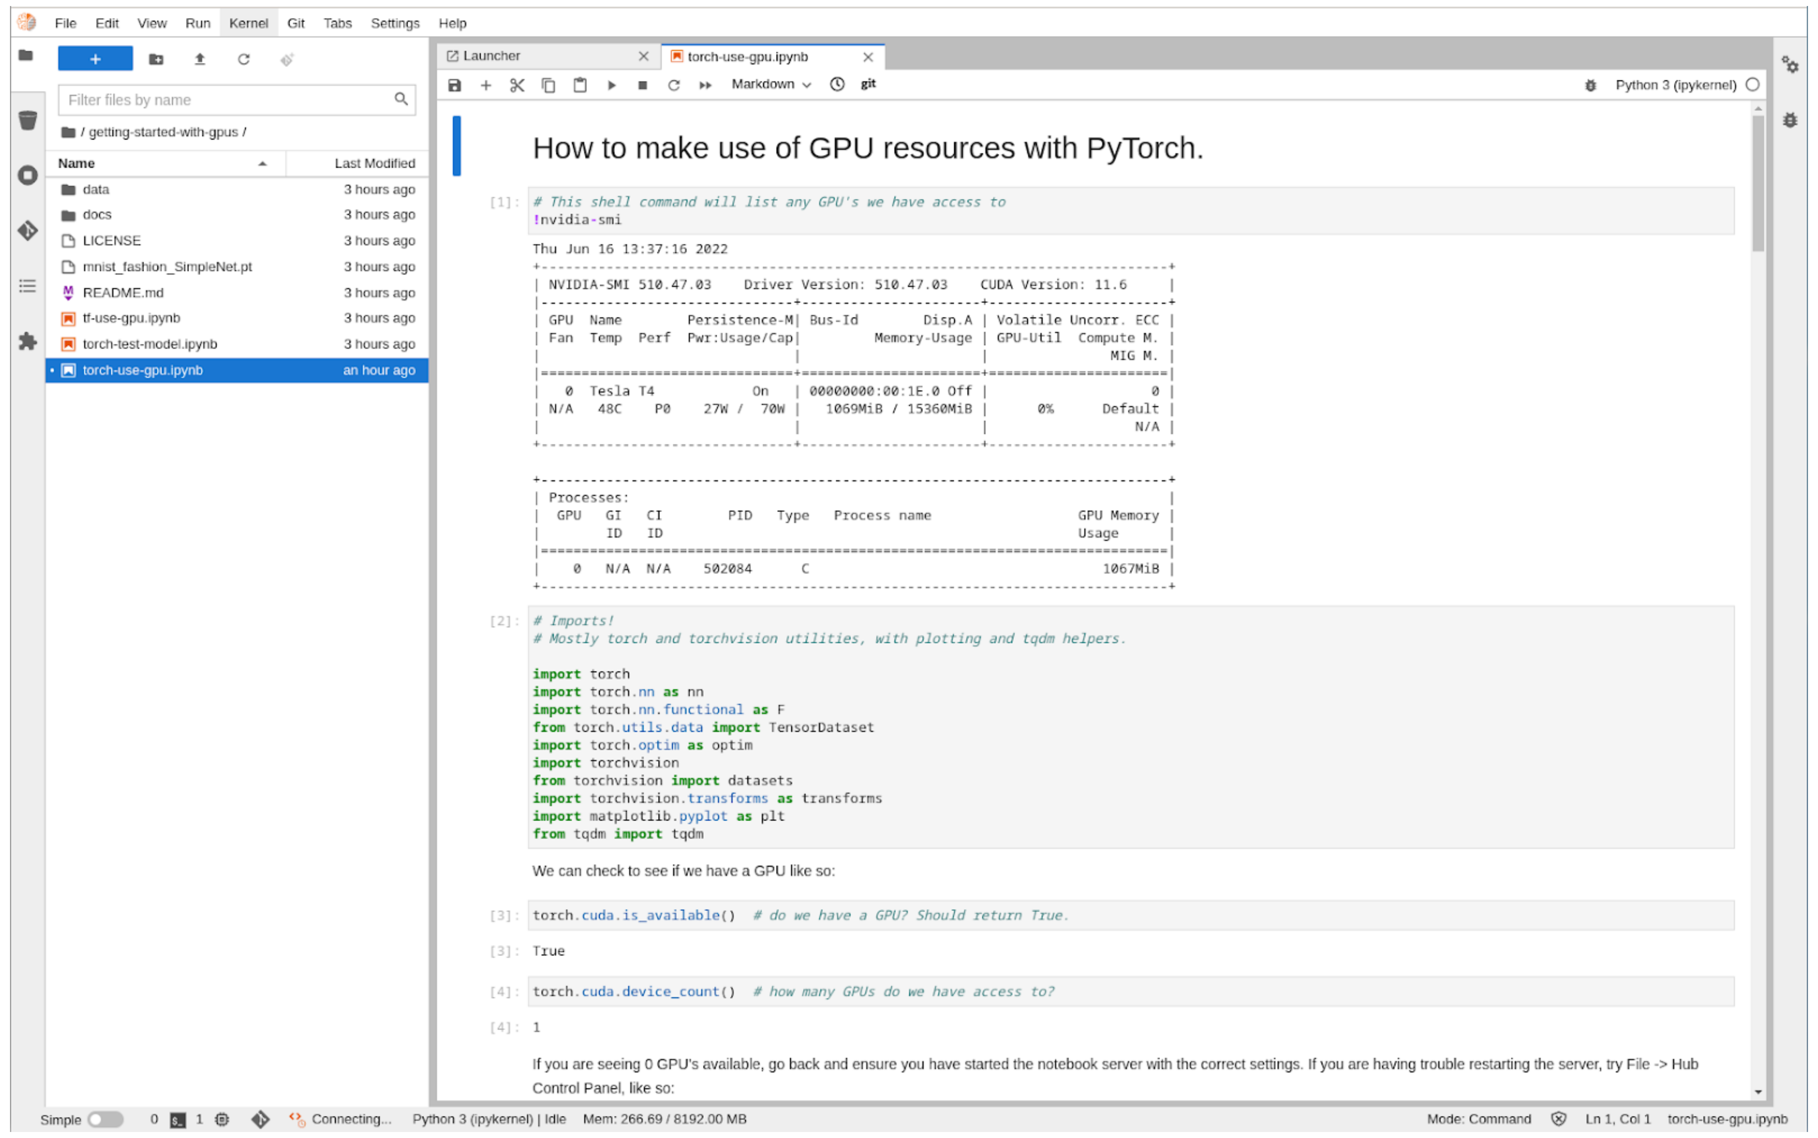

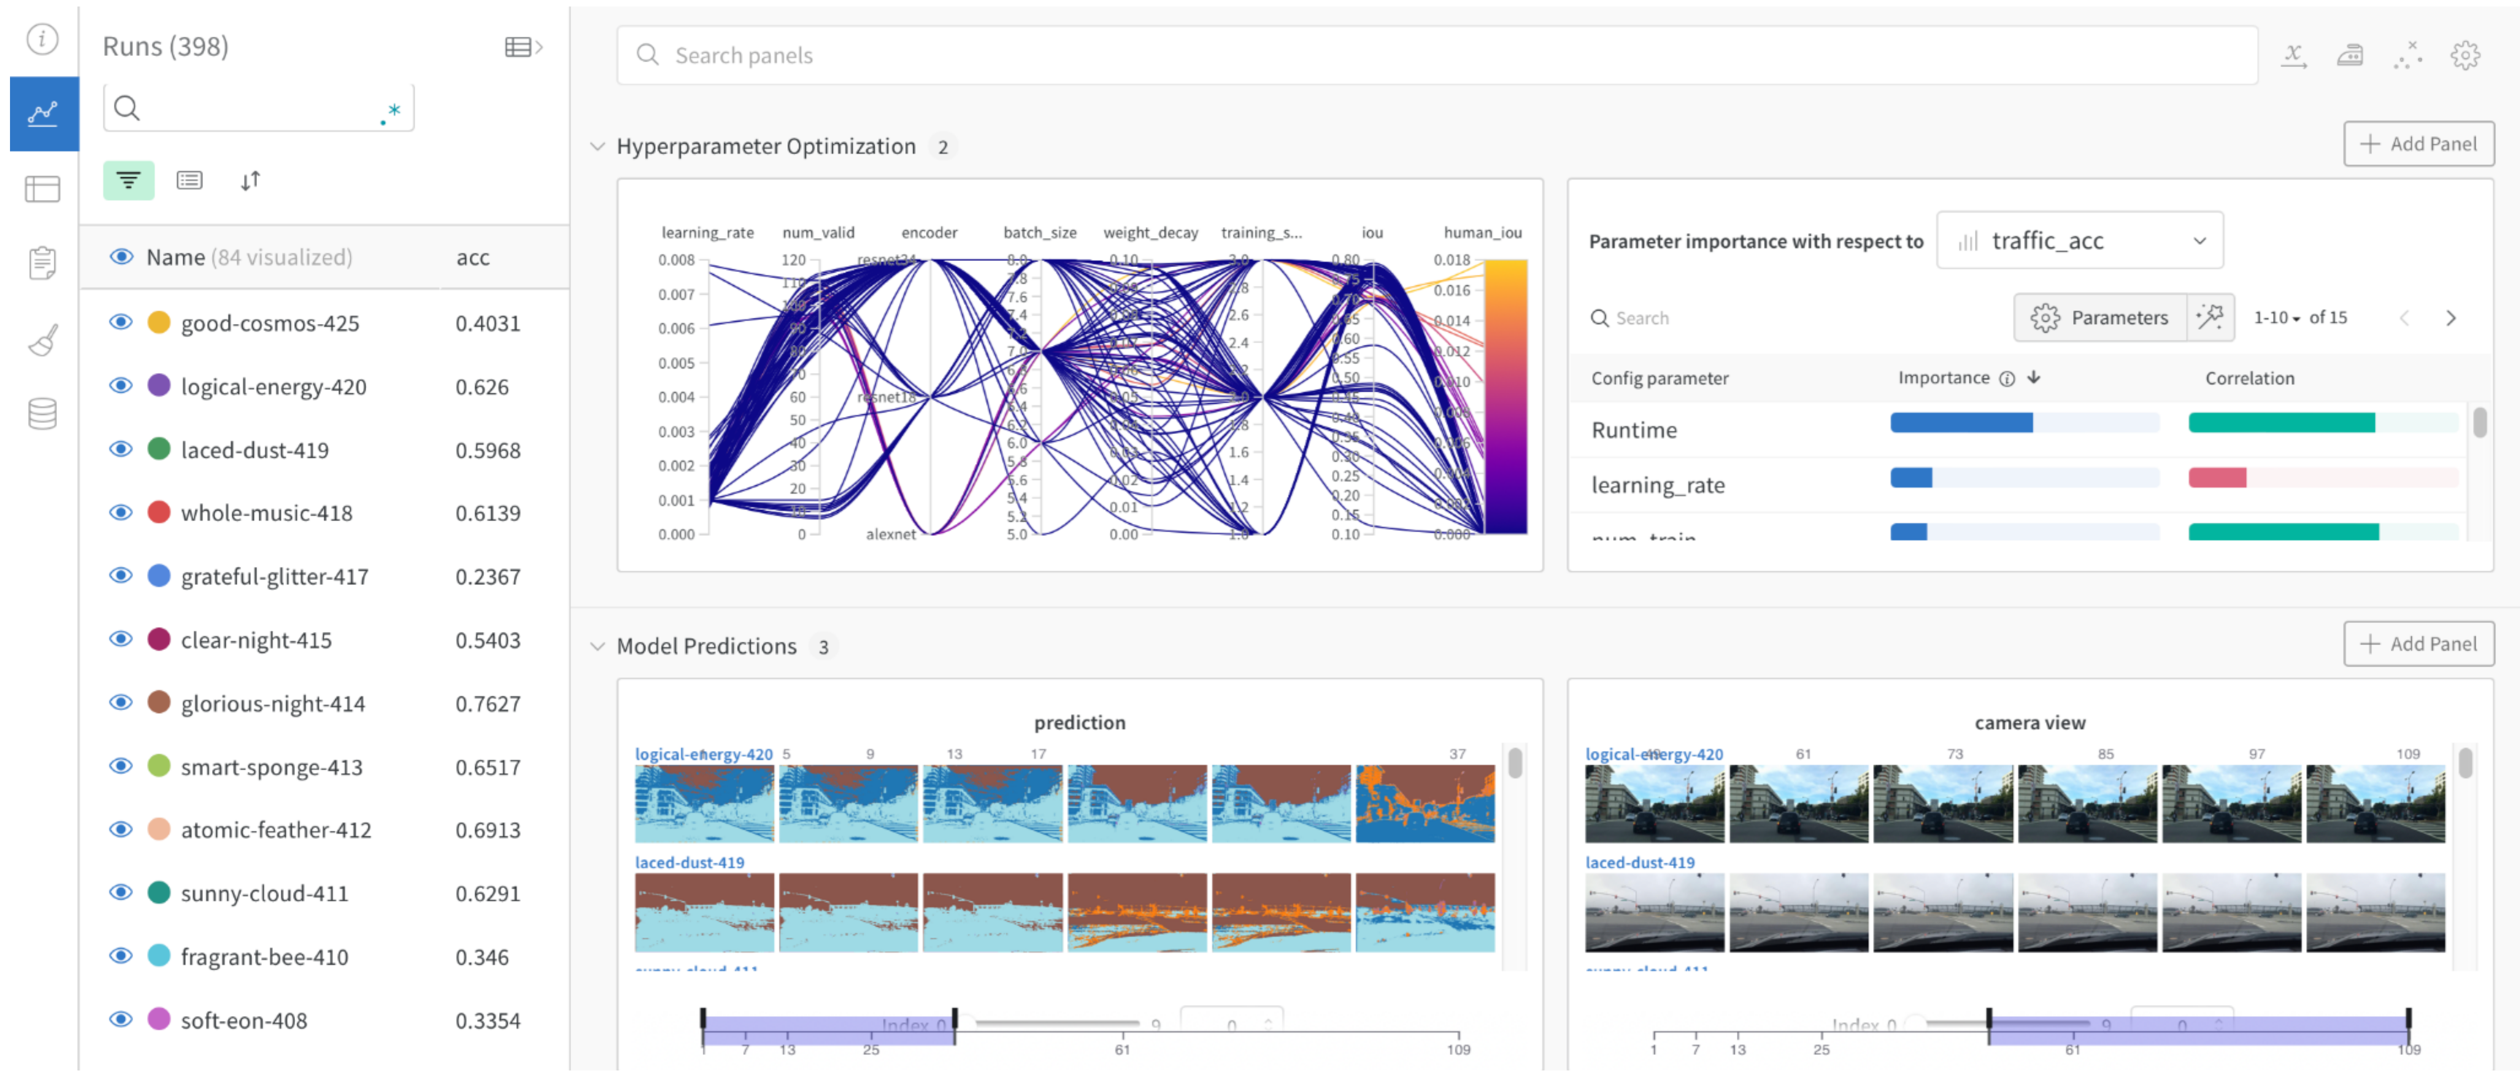

Examples: GPU on a Grid is not enough

ATLAS added GPUs to Grid in ~ 2018 with a clear use case But the ML workflow is different enough that things don't translate directly. And people leave the grid very soon.

GPUs are plentiful and scarce at the same time

Didn't achieve yet the "democratized" access as for CPU

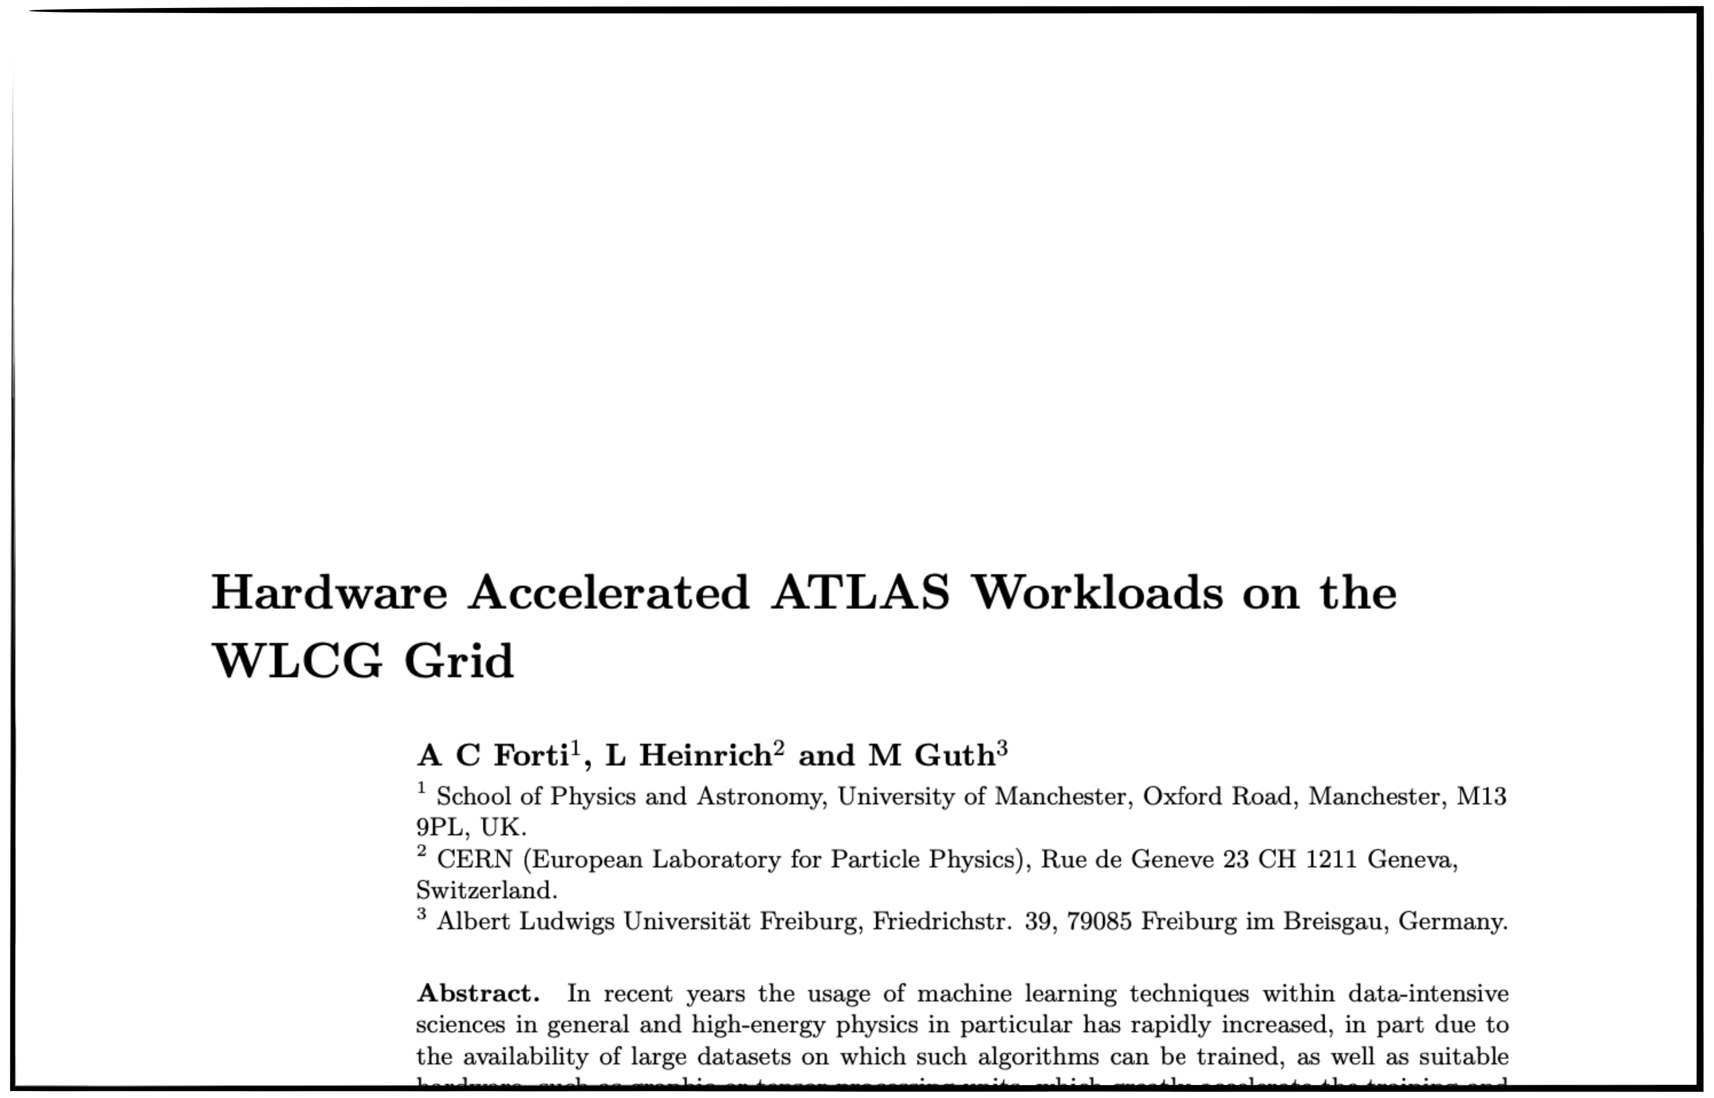

Hardware Accelerated ATLAS Workloads on the WLCG Grid

A C Forti[1], L Heinrich[2] and M Guth[3]

[1] School of Physics and Astronomy, University of Manchester, Oxford Road, Manchester, M13 9PL, UK.

[2] CERN (European Laboratory for Particle Physics), Rue de Geneve 23 CH 1211 Geneva, Switzerland.

[3] Albert Ludwigs Universität Freiburg, Friedrichstr. 39, 79085 Freiburg im Breisgau, Germany.

Abstract. In recent years the usage of machine learning techniques within data-intensive sciences in general and high-energy physics in particular has rapidly increased, in part due to the availability of large datasets on which such algorithms can be trained, as well as suitable

ACAT 2019

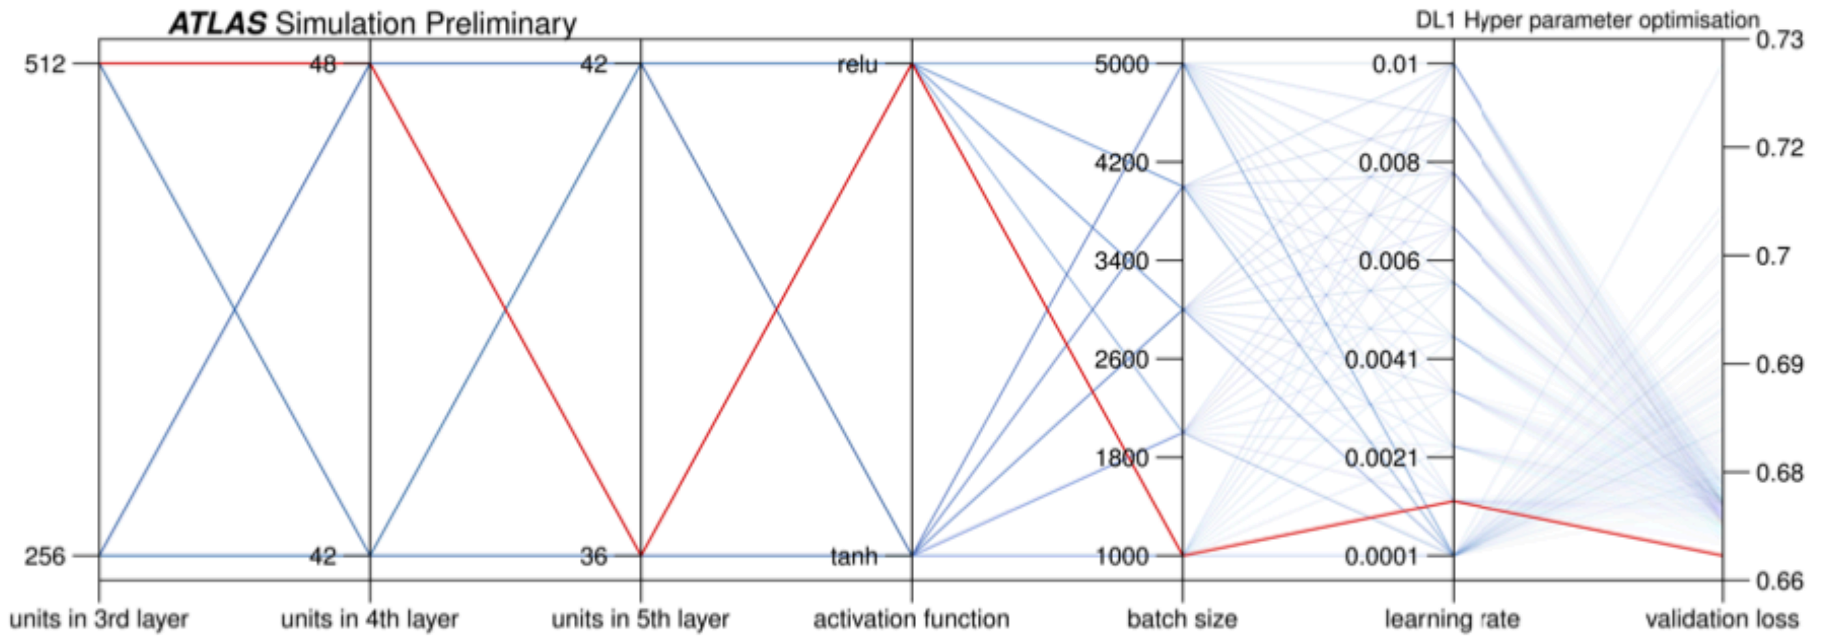

Figure 3. Parallel coordinates plot for 800 different Hyper Parameter combinations. The lines show different combinations of configurations represented in each axis. The last axis shows the neural network loss in the validation sample for a given configuration. The red line shows the Hyper Parameter configuration with the smallest validation loss [7].

Hyperparameter Scan of Flavor Tagging

Analysis Workflows

Analysis aren’t single-use but become useful tools in an of itself.

→ fully-containerized workflows

A Sustainability Argument: Make the most of the data with as little resources as possible

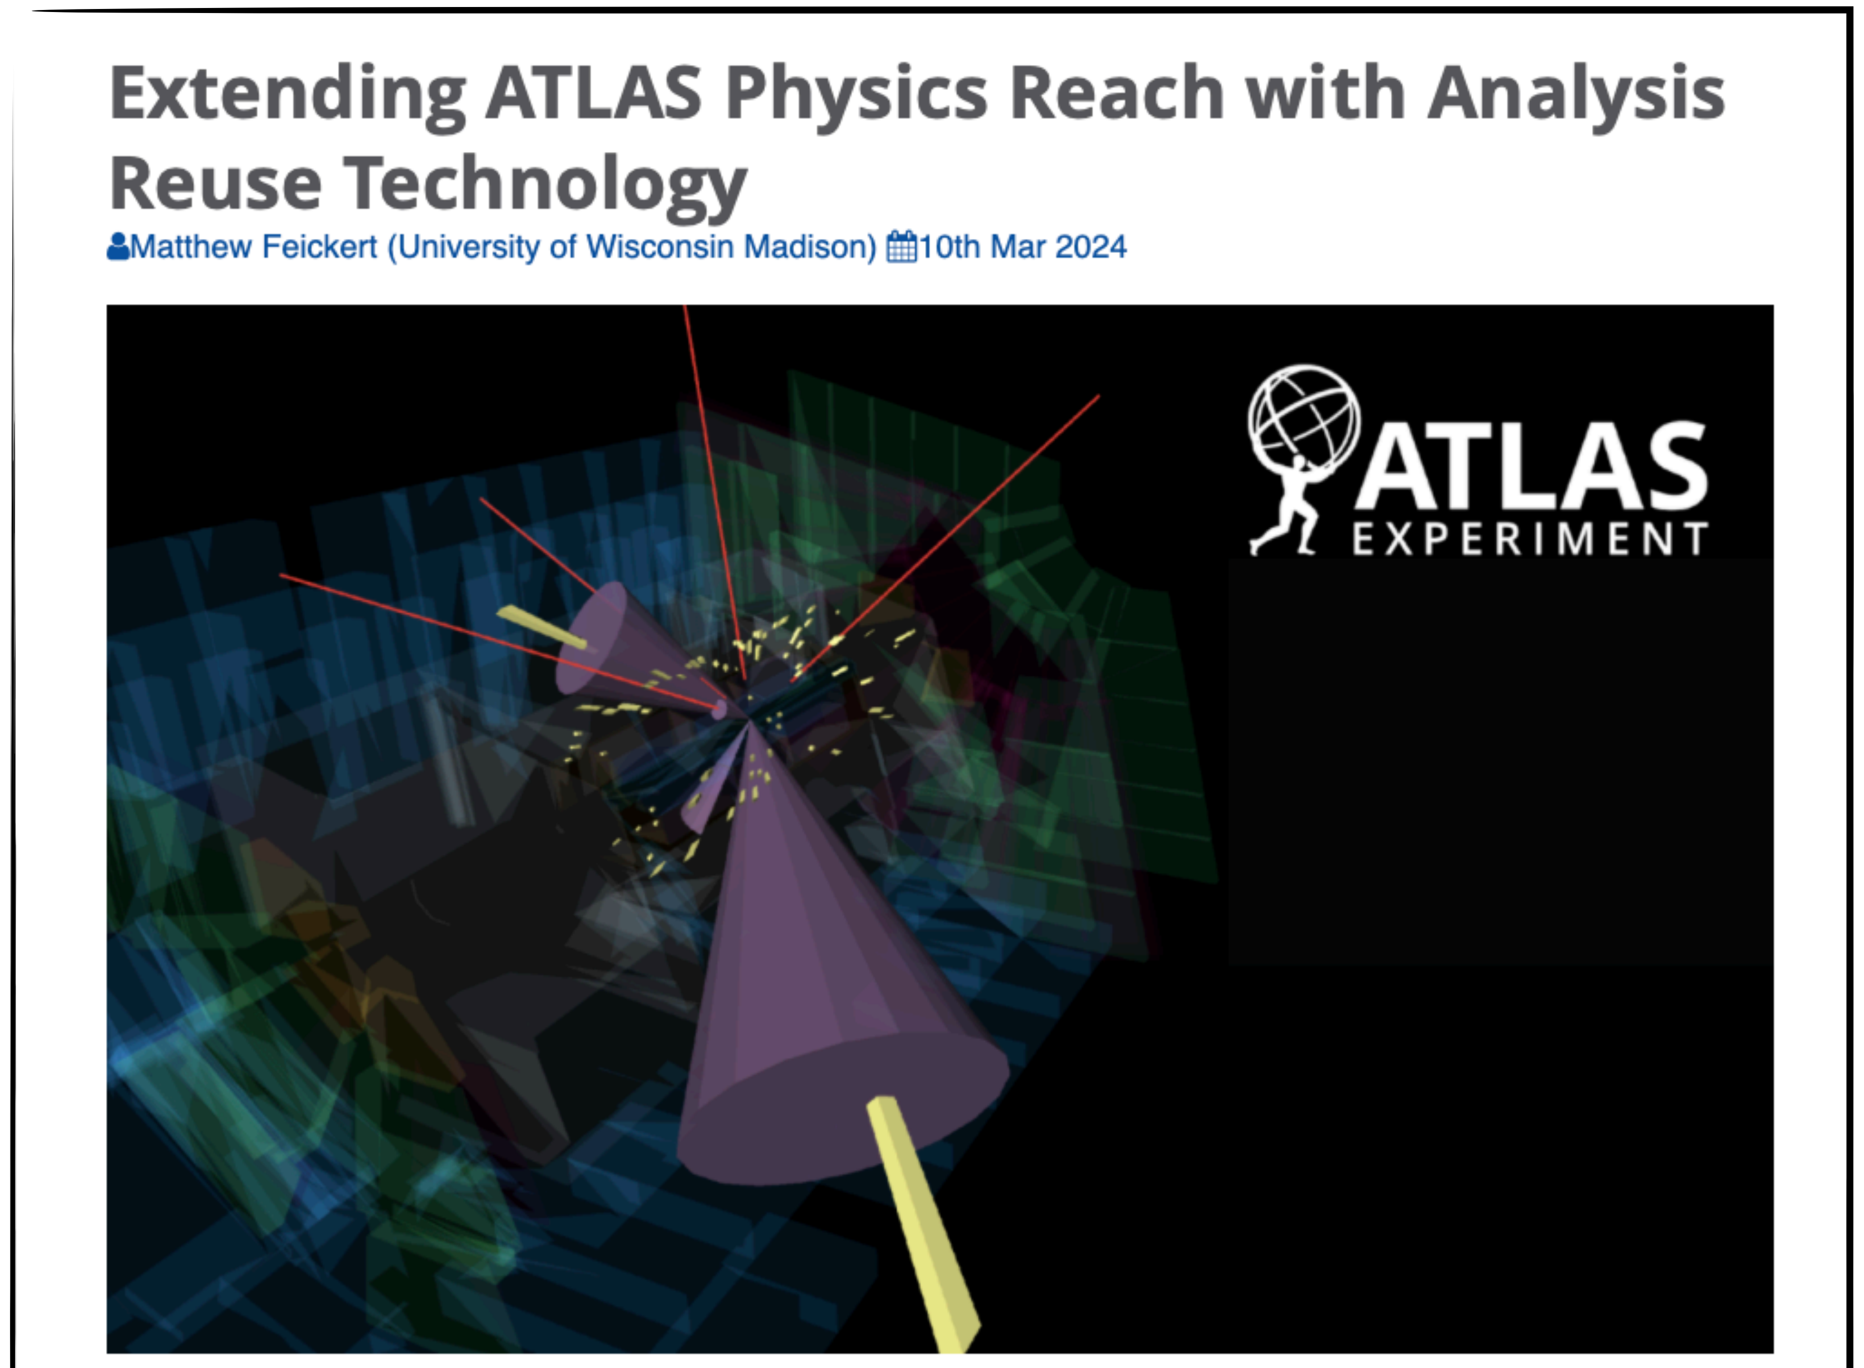

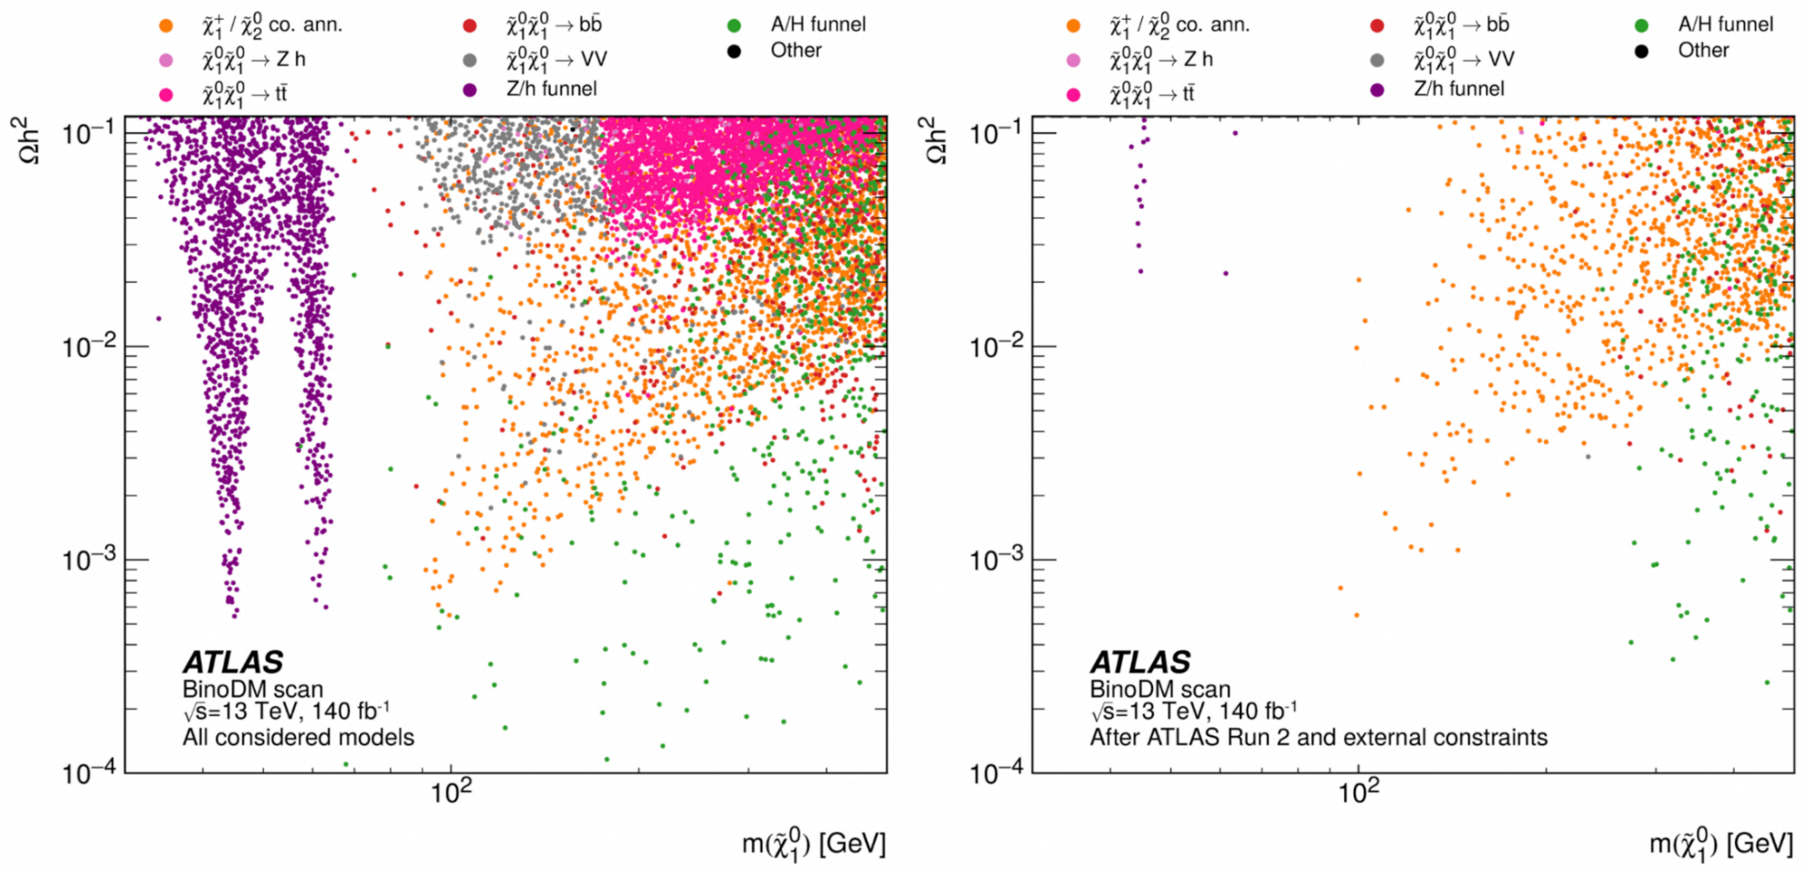

tens of thousands of BSM models analyzed by a single PhD student

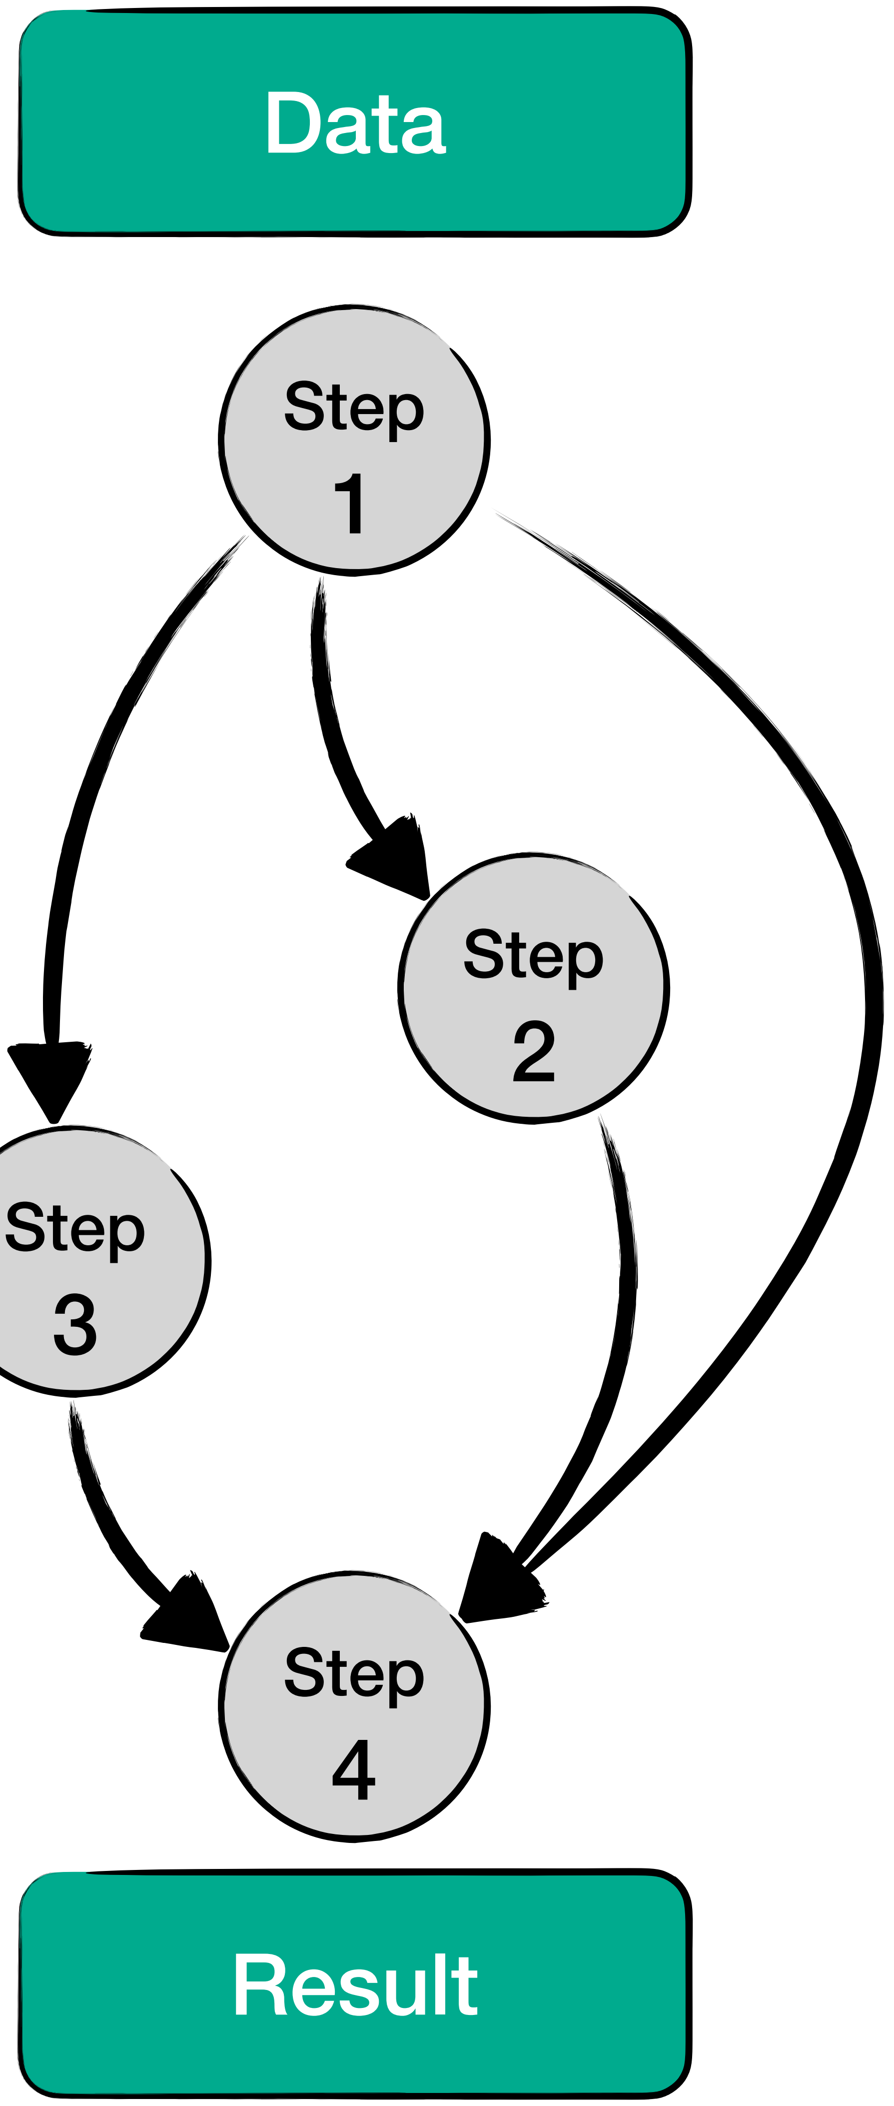

So: want Analysis Infrastructure to do more

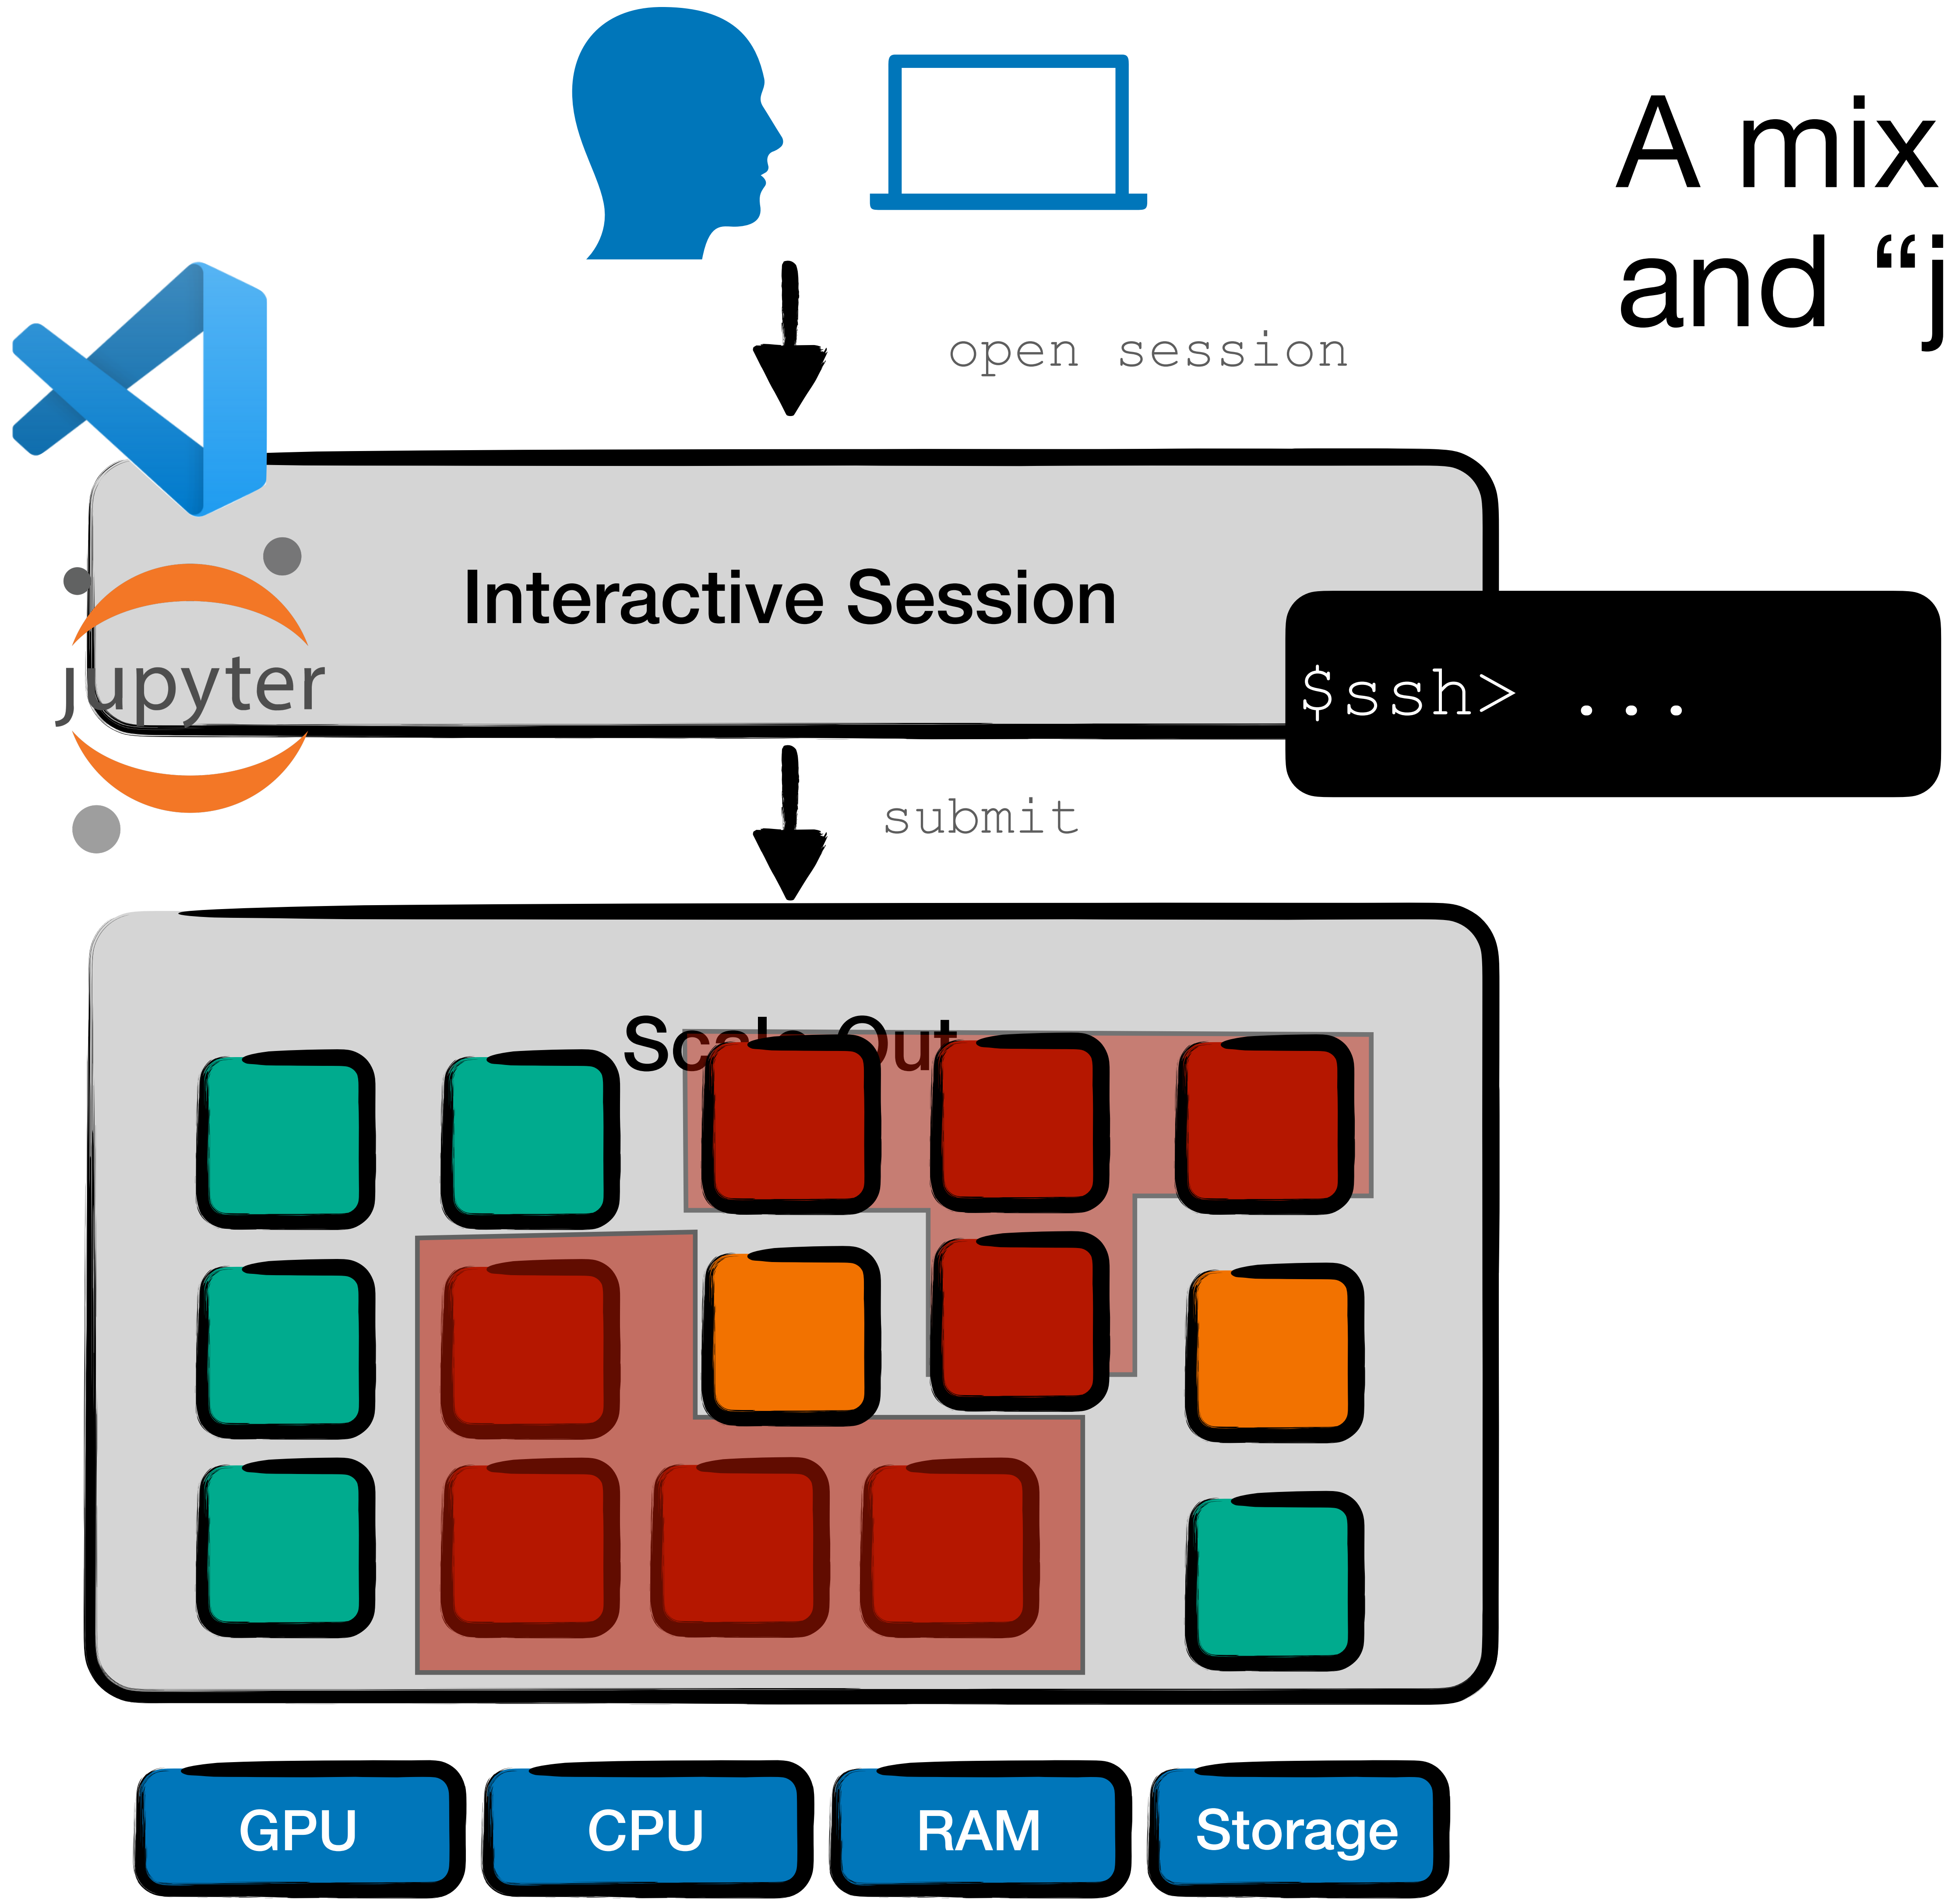

A mix of “services-like” (e.g. JupyterHub) and “job”-like workloads

Infrastructure

It's very noticable that a lot of the AF R&D integrates modern cloud computing tools, especially **Kubernetes**

A lot of the diagrams kind of look similar. Jupyter Hubs, Dask, Batch, ….

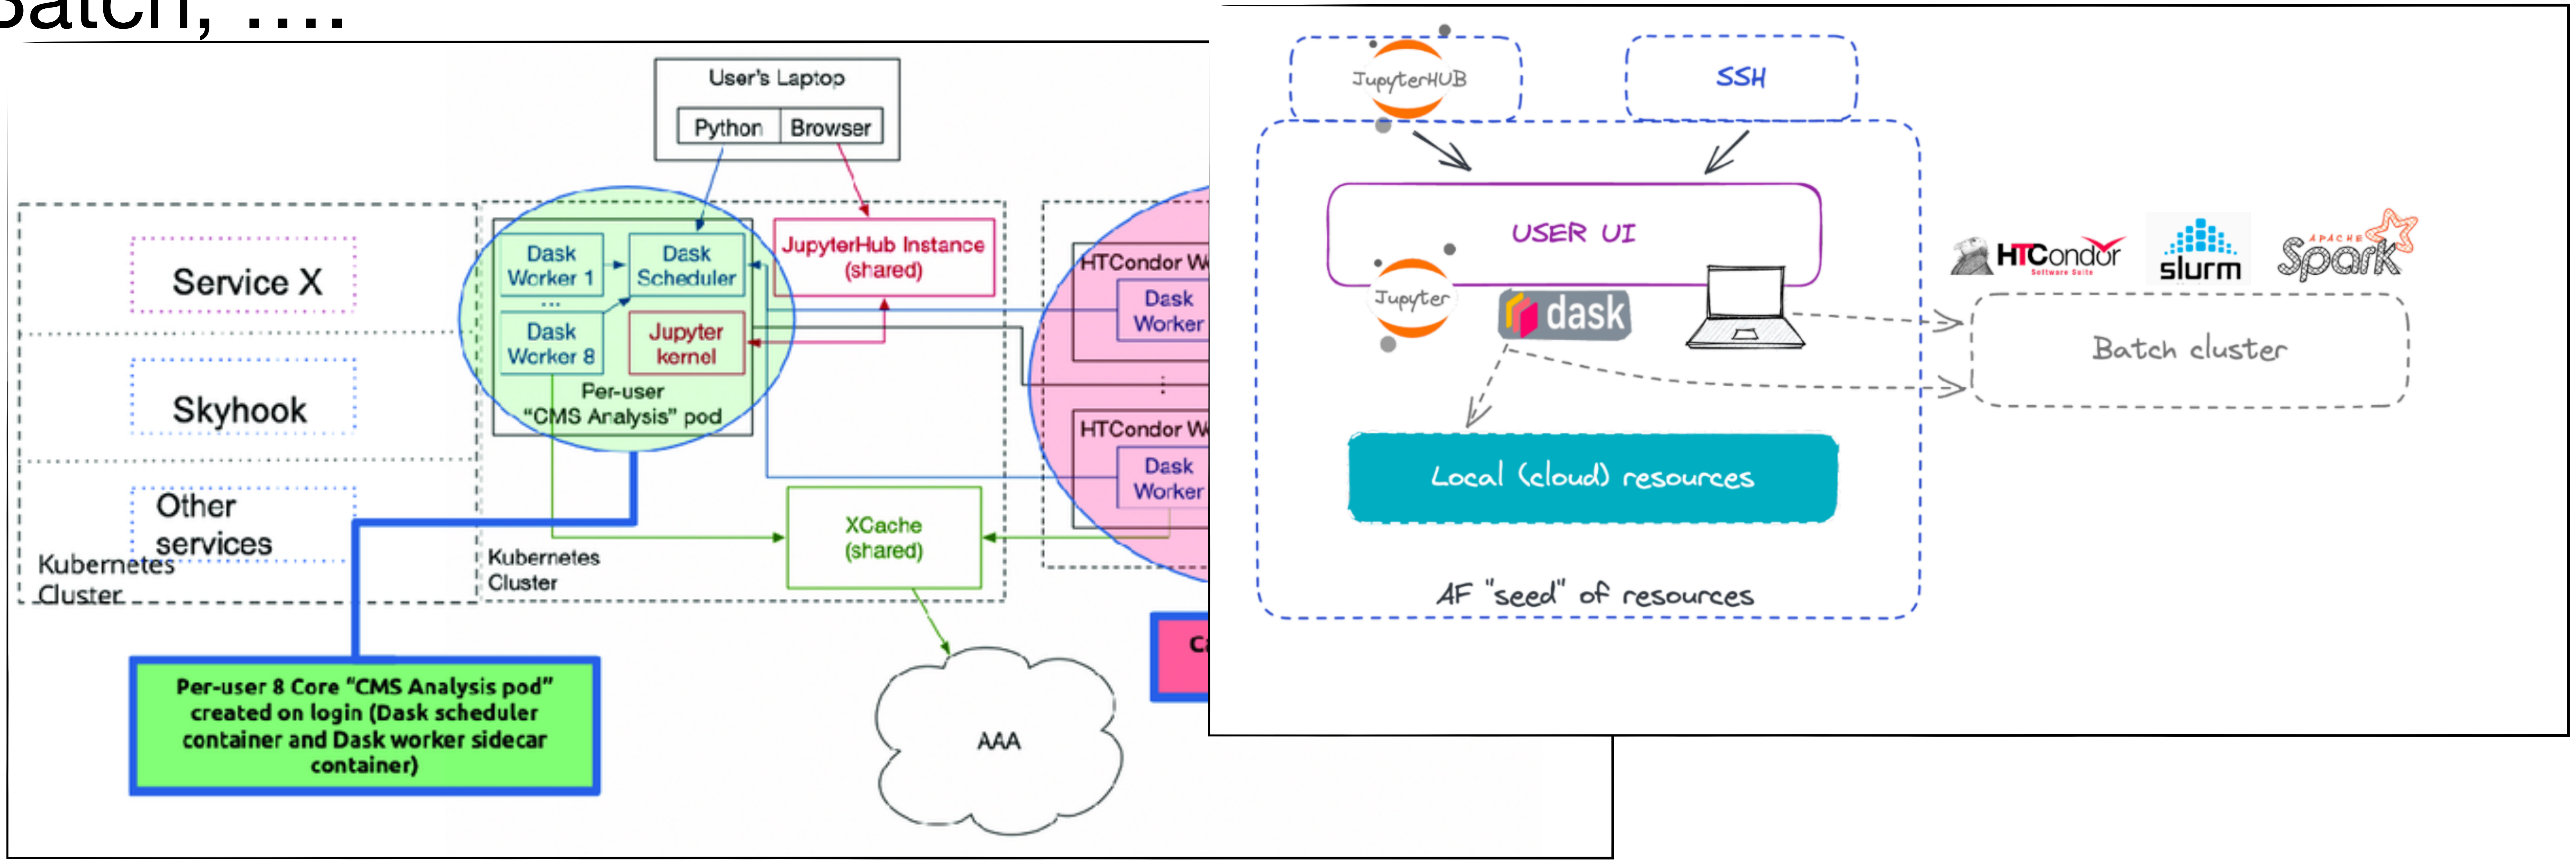

Kubernetes - Lingua Franca

IMHO: fair to say that in a green-field environment, Kubernetes would be a **very strong contender for a infrastructure foundation** / fabric.
→ designed for mixed workloads + dynamic scaling. Scales very well.
→ increasingly what tools we might want to use build towards

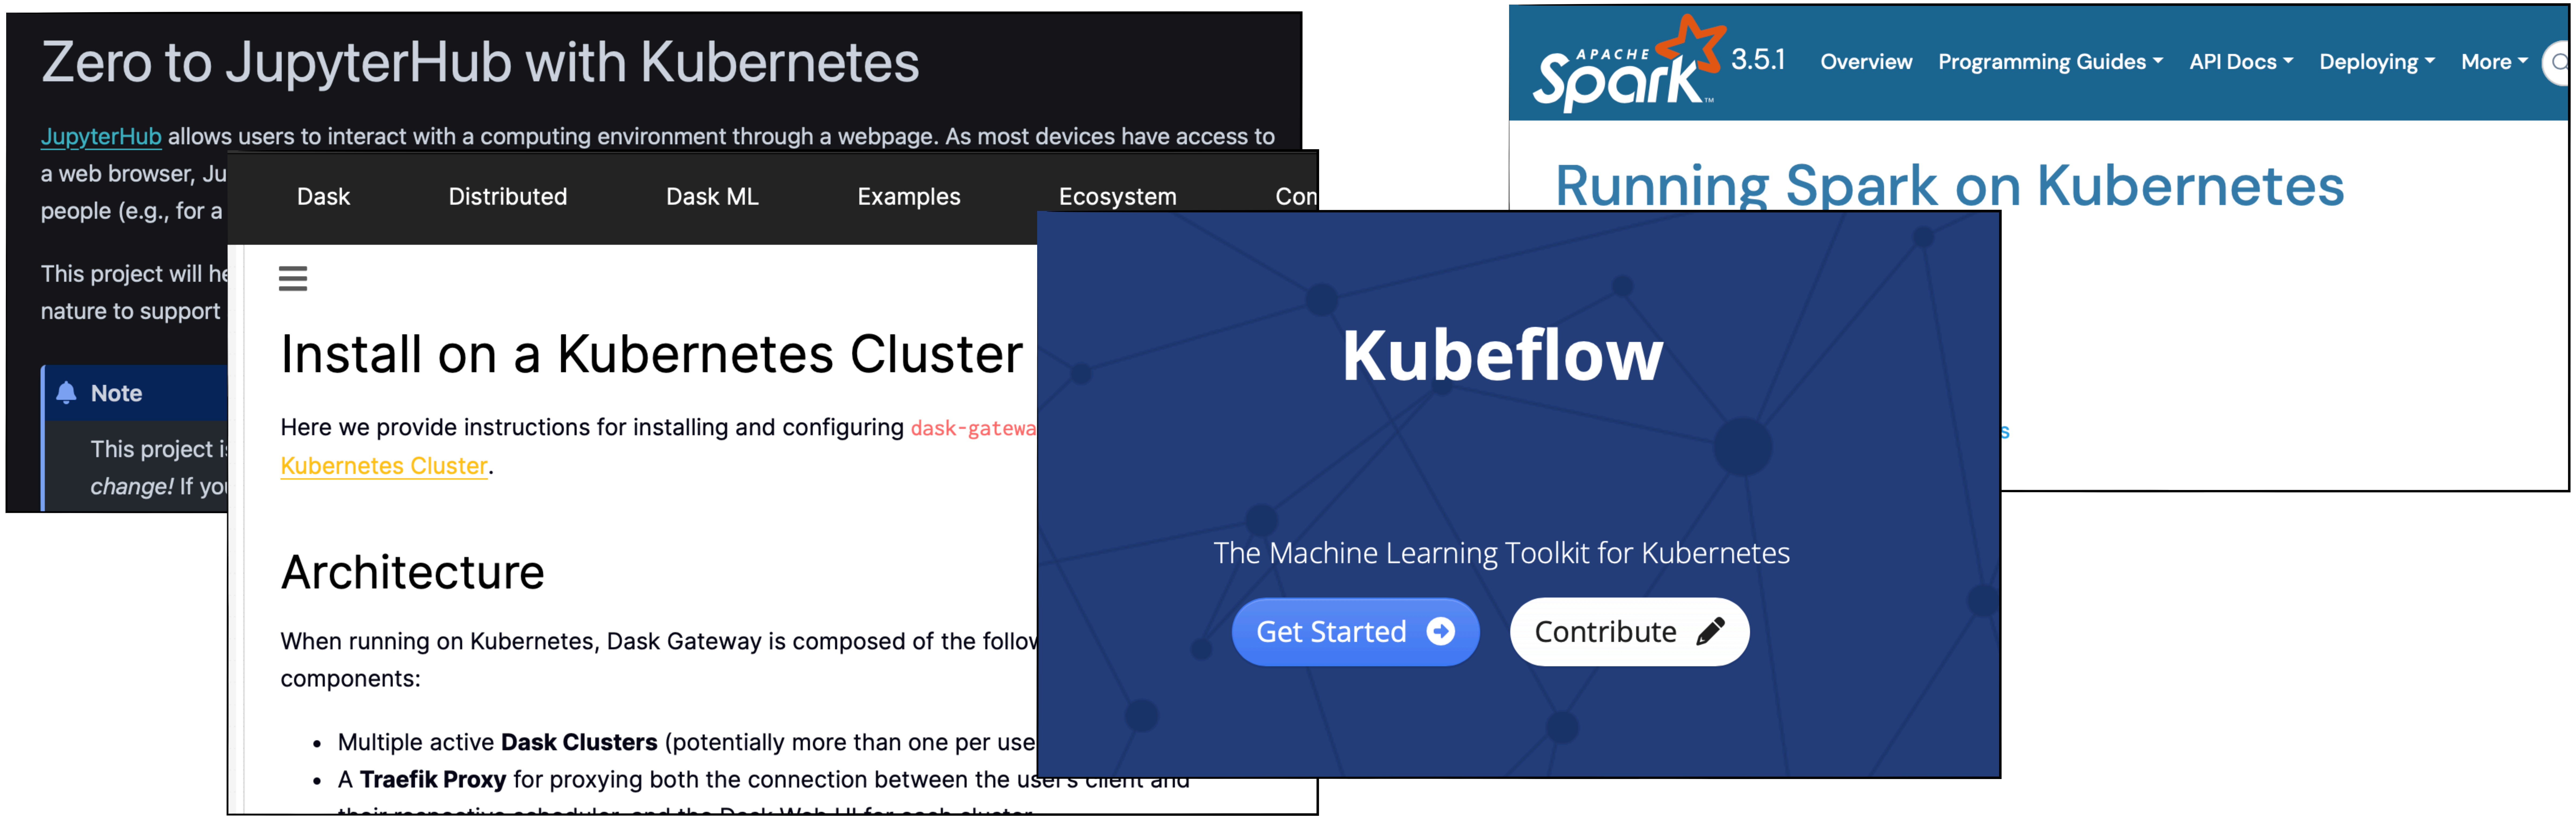

Kubernetes - Lingua Franca

IMHO: fair to say that in a green-field environment, Kubernetes would be a **very strong contender for a infrastructure foundation** / fabric.

→“k8s in Academia” still a bit unsolved (but e.g. CERN a big success)

→ perfect research project e.g. ErUM Data? (also k8s + Storage…)

→ concise way to communicate e.g. req’s to HPC (HPC as Cloud Provider)

→ collaboration pathway with industry + important skillset for our people

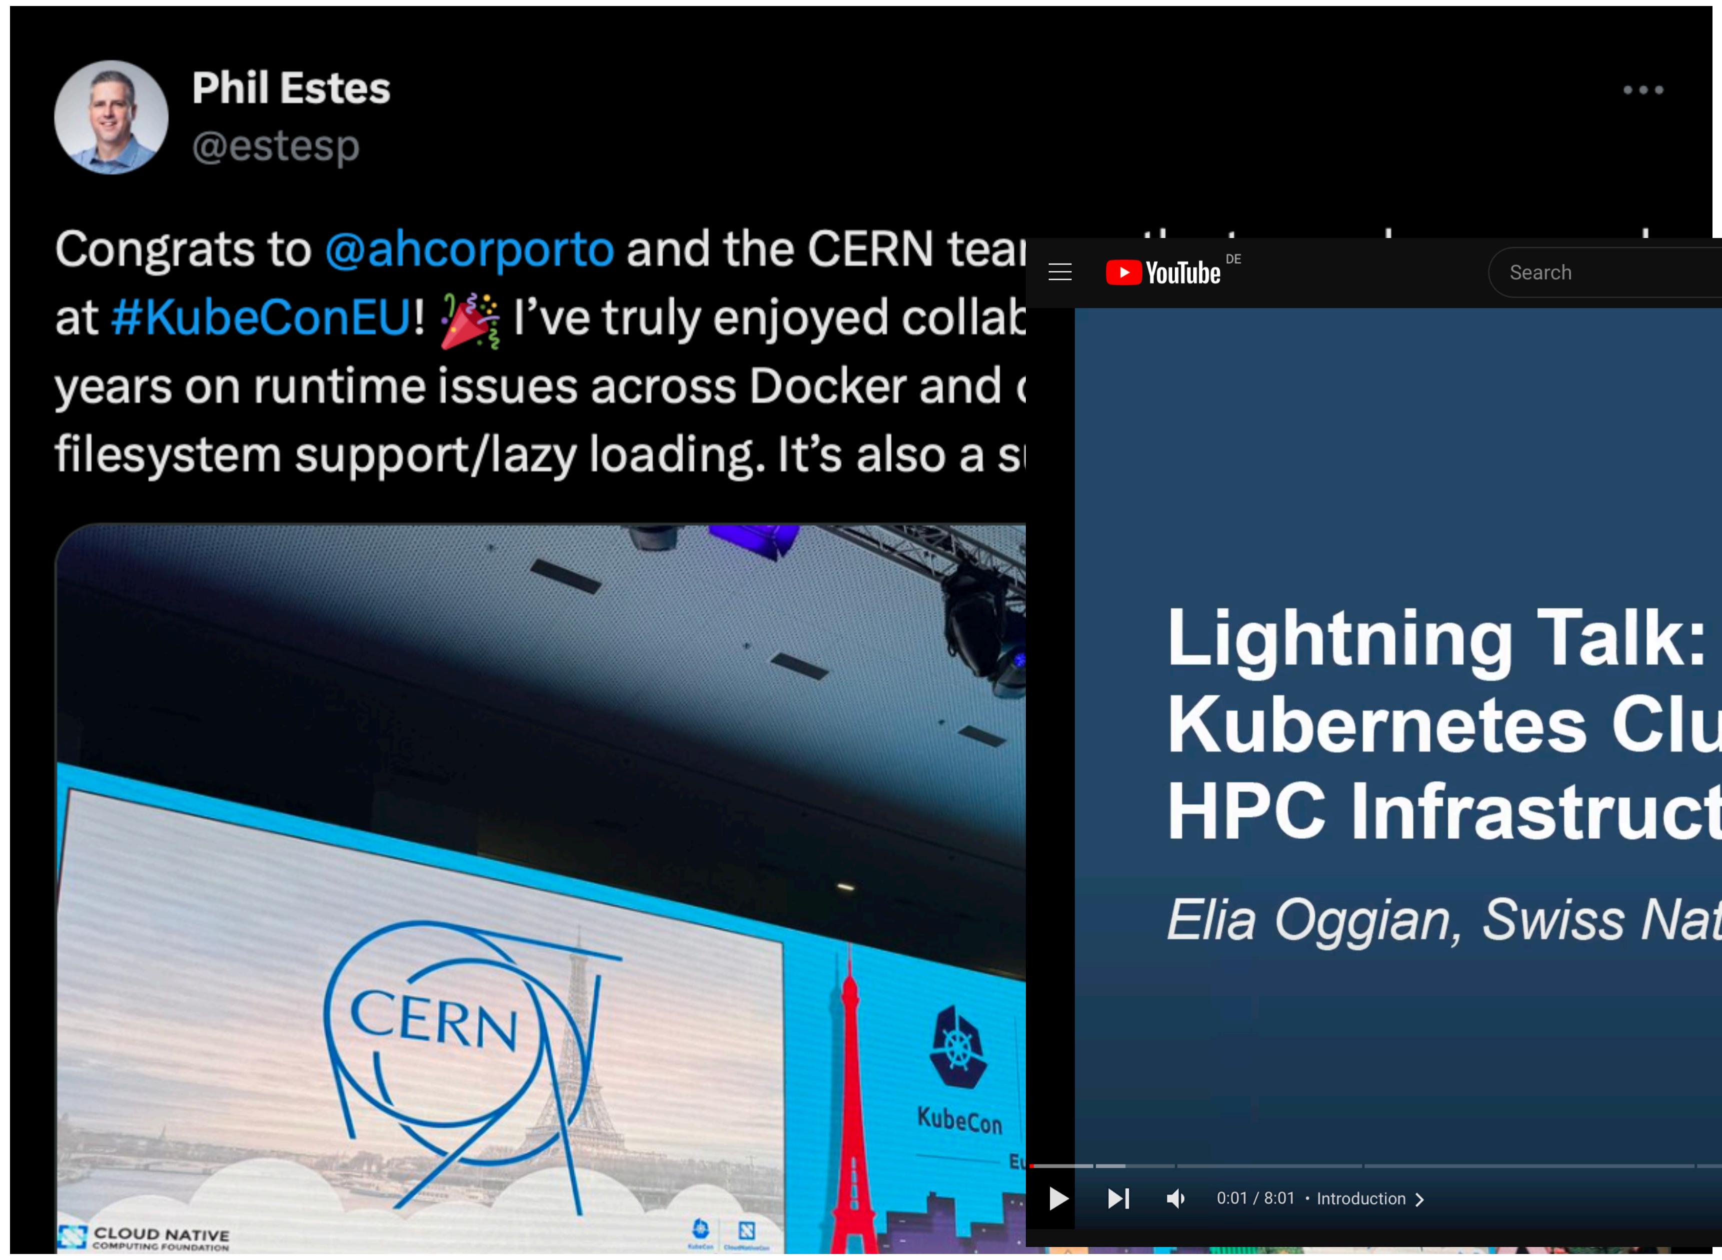

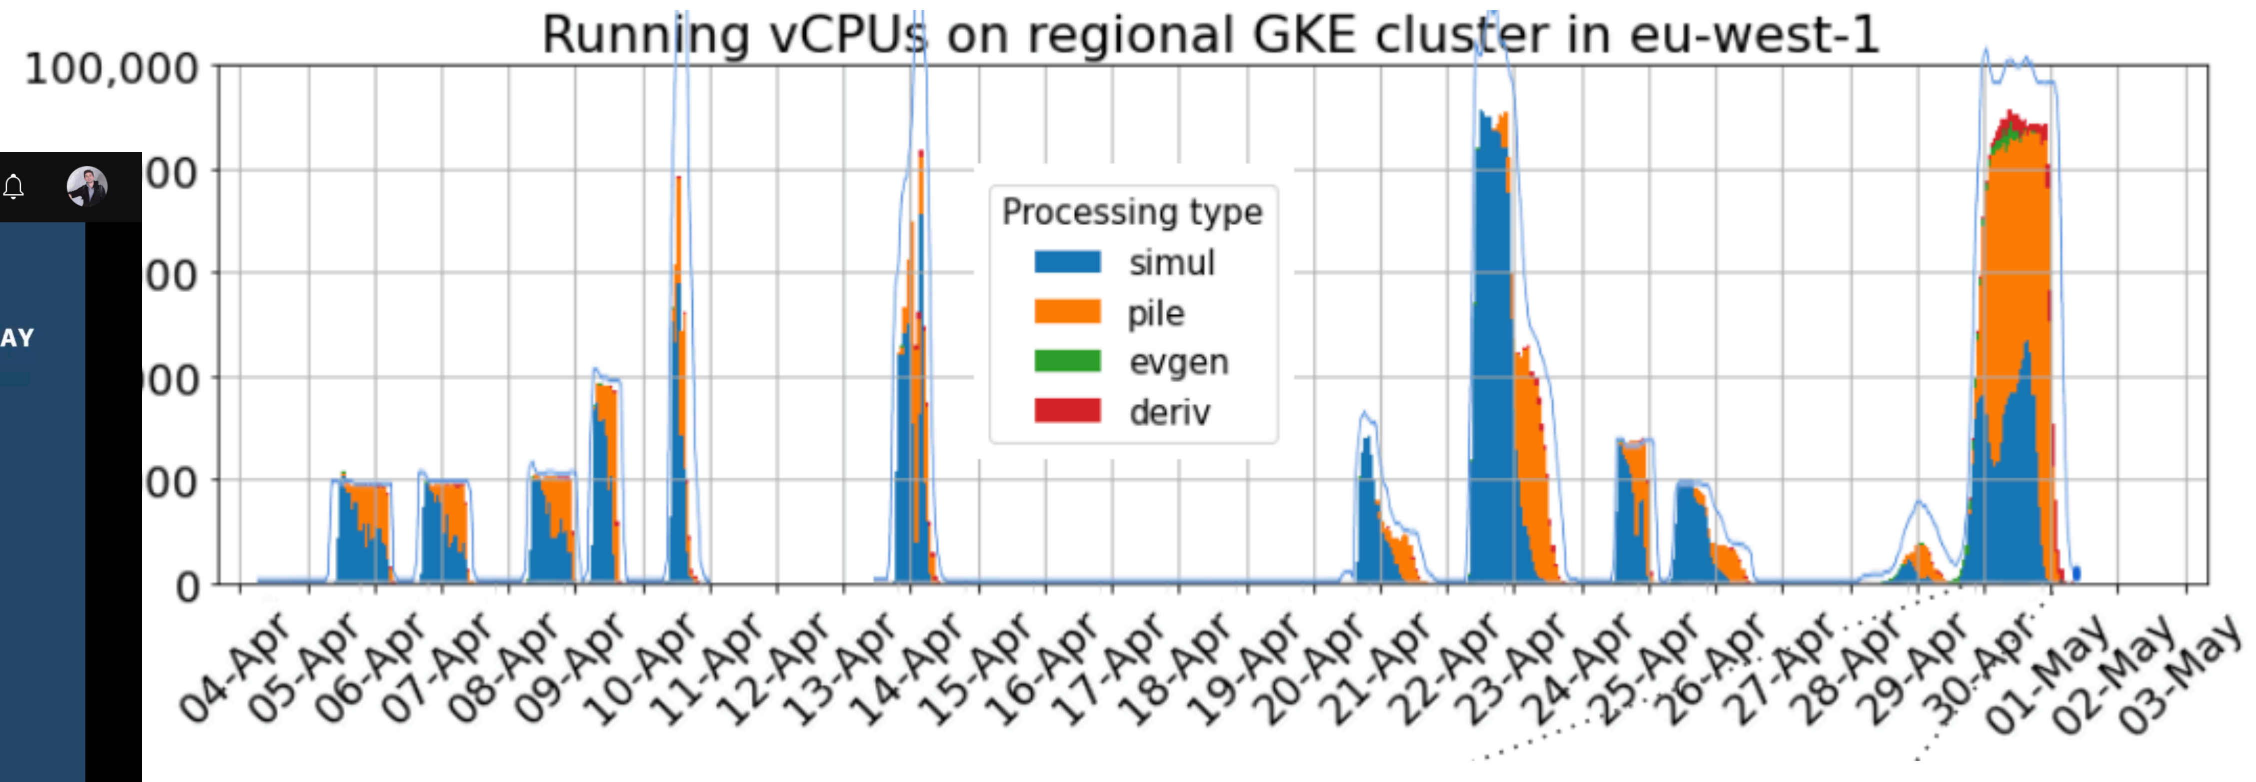

Scale out to 100k cores with a single k8s cluster + Grid

More Details - HSF AF Whitepaper

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 [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], José Hernández [12], Brian Bockelman [25], Eric Lancon [9], James Frost [26], Andrea Sciaba [6], Markus Schulz [6], Caterina Doglioni [2], Guenter Duckeck [27], Tibor Šimko [6], Luke Kreczko [28], Farid Ould-Saada [29], Eirik Gramstad [29], James Catmore [29], Matthew Feickert [15], Iason Krommydas [30], Shawn McKee [31], Tomas Lindén [32], Jim Pivarski [33], Peter Elmer [33], Ianna Osborne [33], Andrzej Novak [34]

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, 26 University of Oxford, 27 Ludwigs-Maximilians Universität München, 28 University of Bristol, 29 University of Oslo, 30 Rice University, 31 University of Michigan, 32 Helsinki Institute of Physics, 33 Princeton University, 34 Massachusetts Institute of Technology

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

Tried to summarize and contextualize developments.
→ evolution / extension of Capabilities rather than something completely new

(Couldn't cover everything)

Open to endorsement authorship until April 1. More German Visibility?
→ useful as demonstrable "Vorarbeit" for ErUM-Data etc

Let me know & I can add you

Where to go from here

Thomas yesterday: Lots of demonstrators, small-scale showcases. Great to explore (let a thousand flowers bloom, see what sticks), **but real deployment is very different.**

Next phase should be an **effort to push these ideas into production** systems. Given these developments, what's the vision for a production AF?

What assumptions break, are untenable, don't scale, etc.

But still be ambitious: Opportunity to improve analysis experience.

End.