The Grid on the Cloud

A Vision of Analysis Facilities on 100k Cores

Lukas Heinrich, TUM

Showcasing work by

F. Megino, Tibor Simko, Nikolai Hartmann, Alessandra Forti, A. Hartin, G. Duckeck, R. Walker
P. Gadow, J. Elmsheuser, A. Klimentor, K De, M. Lassnig and many more

ATLAS / LHC centric, but hopefully with wider applicability

LHC Computing is a big success

… after all it gave us this 10 years go …

It would have been impossible to release physics results so quickly without the outstanding performance of the Grid (including the CERN Tier-0)

Number of concurrent ATLAS jobs Jan-July 2012

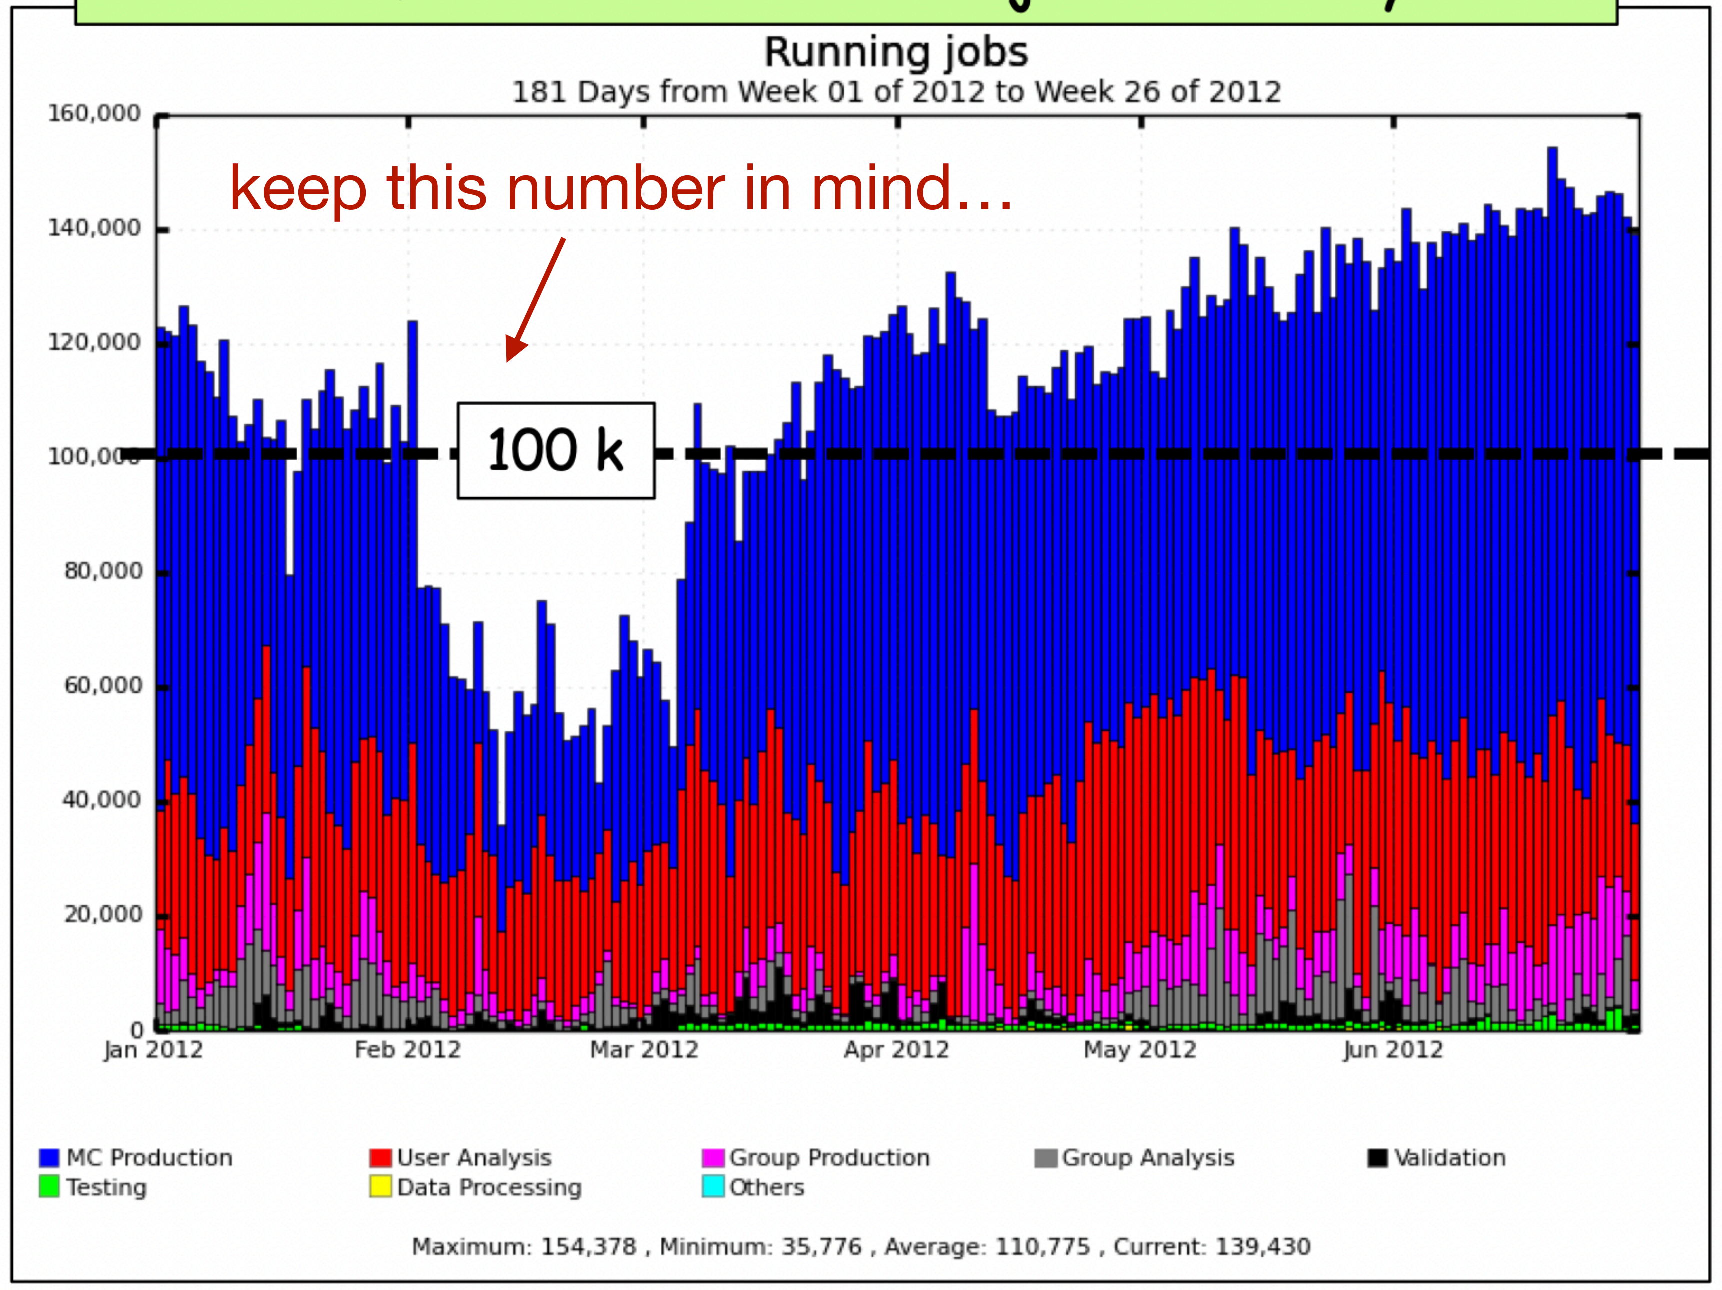

Includes MC production, user and group analysis at CERN, 10 Tier1-s, ~ 70 Tier-2 federations → > 80 sites

> 1500 distinct ATLAS users do analysis on the GRID

- Available resources fully used/stressed (beyond pledges in some cases)
- Massive production of 8 TeV Monte Carlo samples
- Very effective and flexible Computing Model and Operation team → accommodate high trigger rates and pile-up, intense MC simulation, analysis demands from worldwide users (through e.g. dynamic data placement)

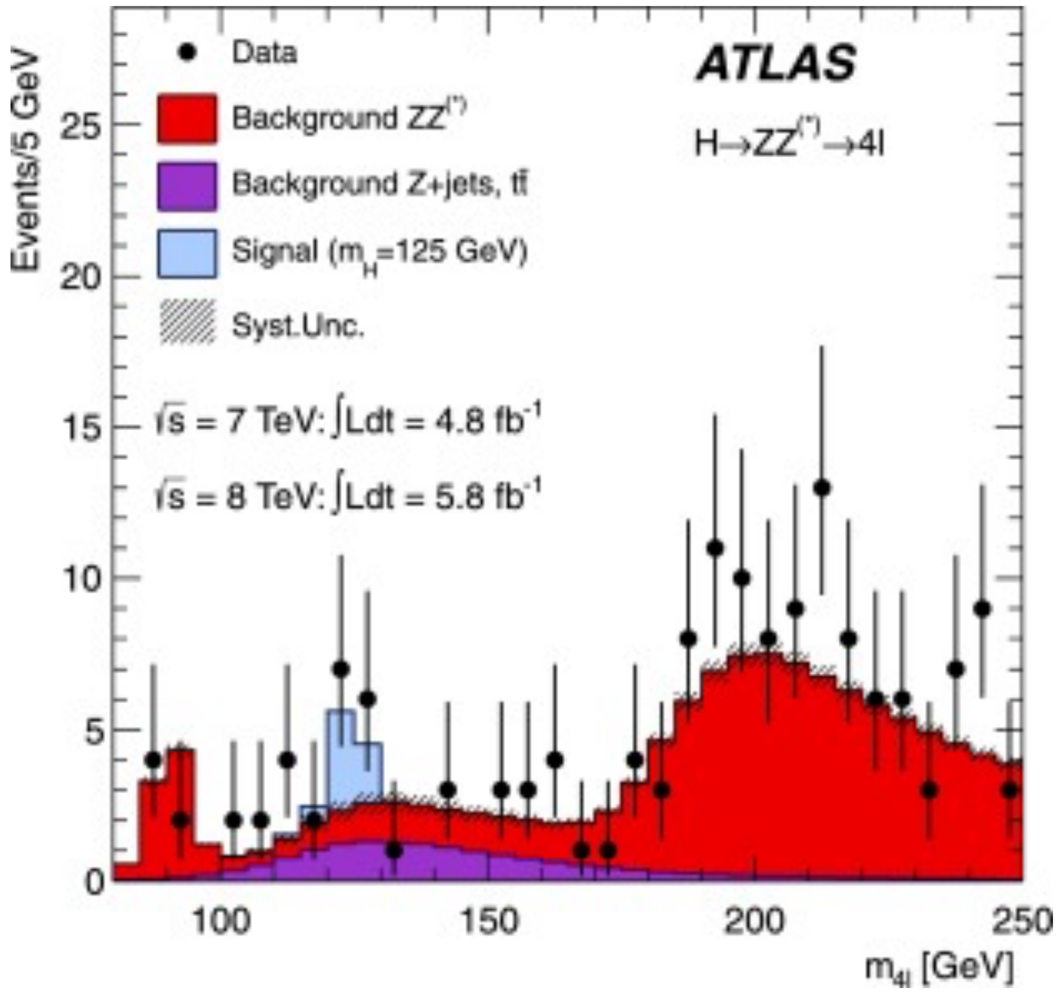

The Worldwide LHC Computing Grid

… the cloud before there was a cloud

- Globally federated computing long before there was a cloud
- Commodity x86 platform
- Linux as a common platform choice
- Tight software control to bridge heterogeneity (CVMFS)
- Focused solely on Batch computing
- Distribute, process and analyse LHC data

ATLAS stats:
165 data centers
40 countries
800k vCPUs
700 PB managed data
3 Tbps network

The Run-3 is here

A glimpse into the future

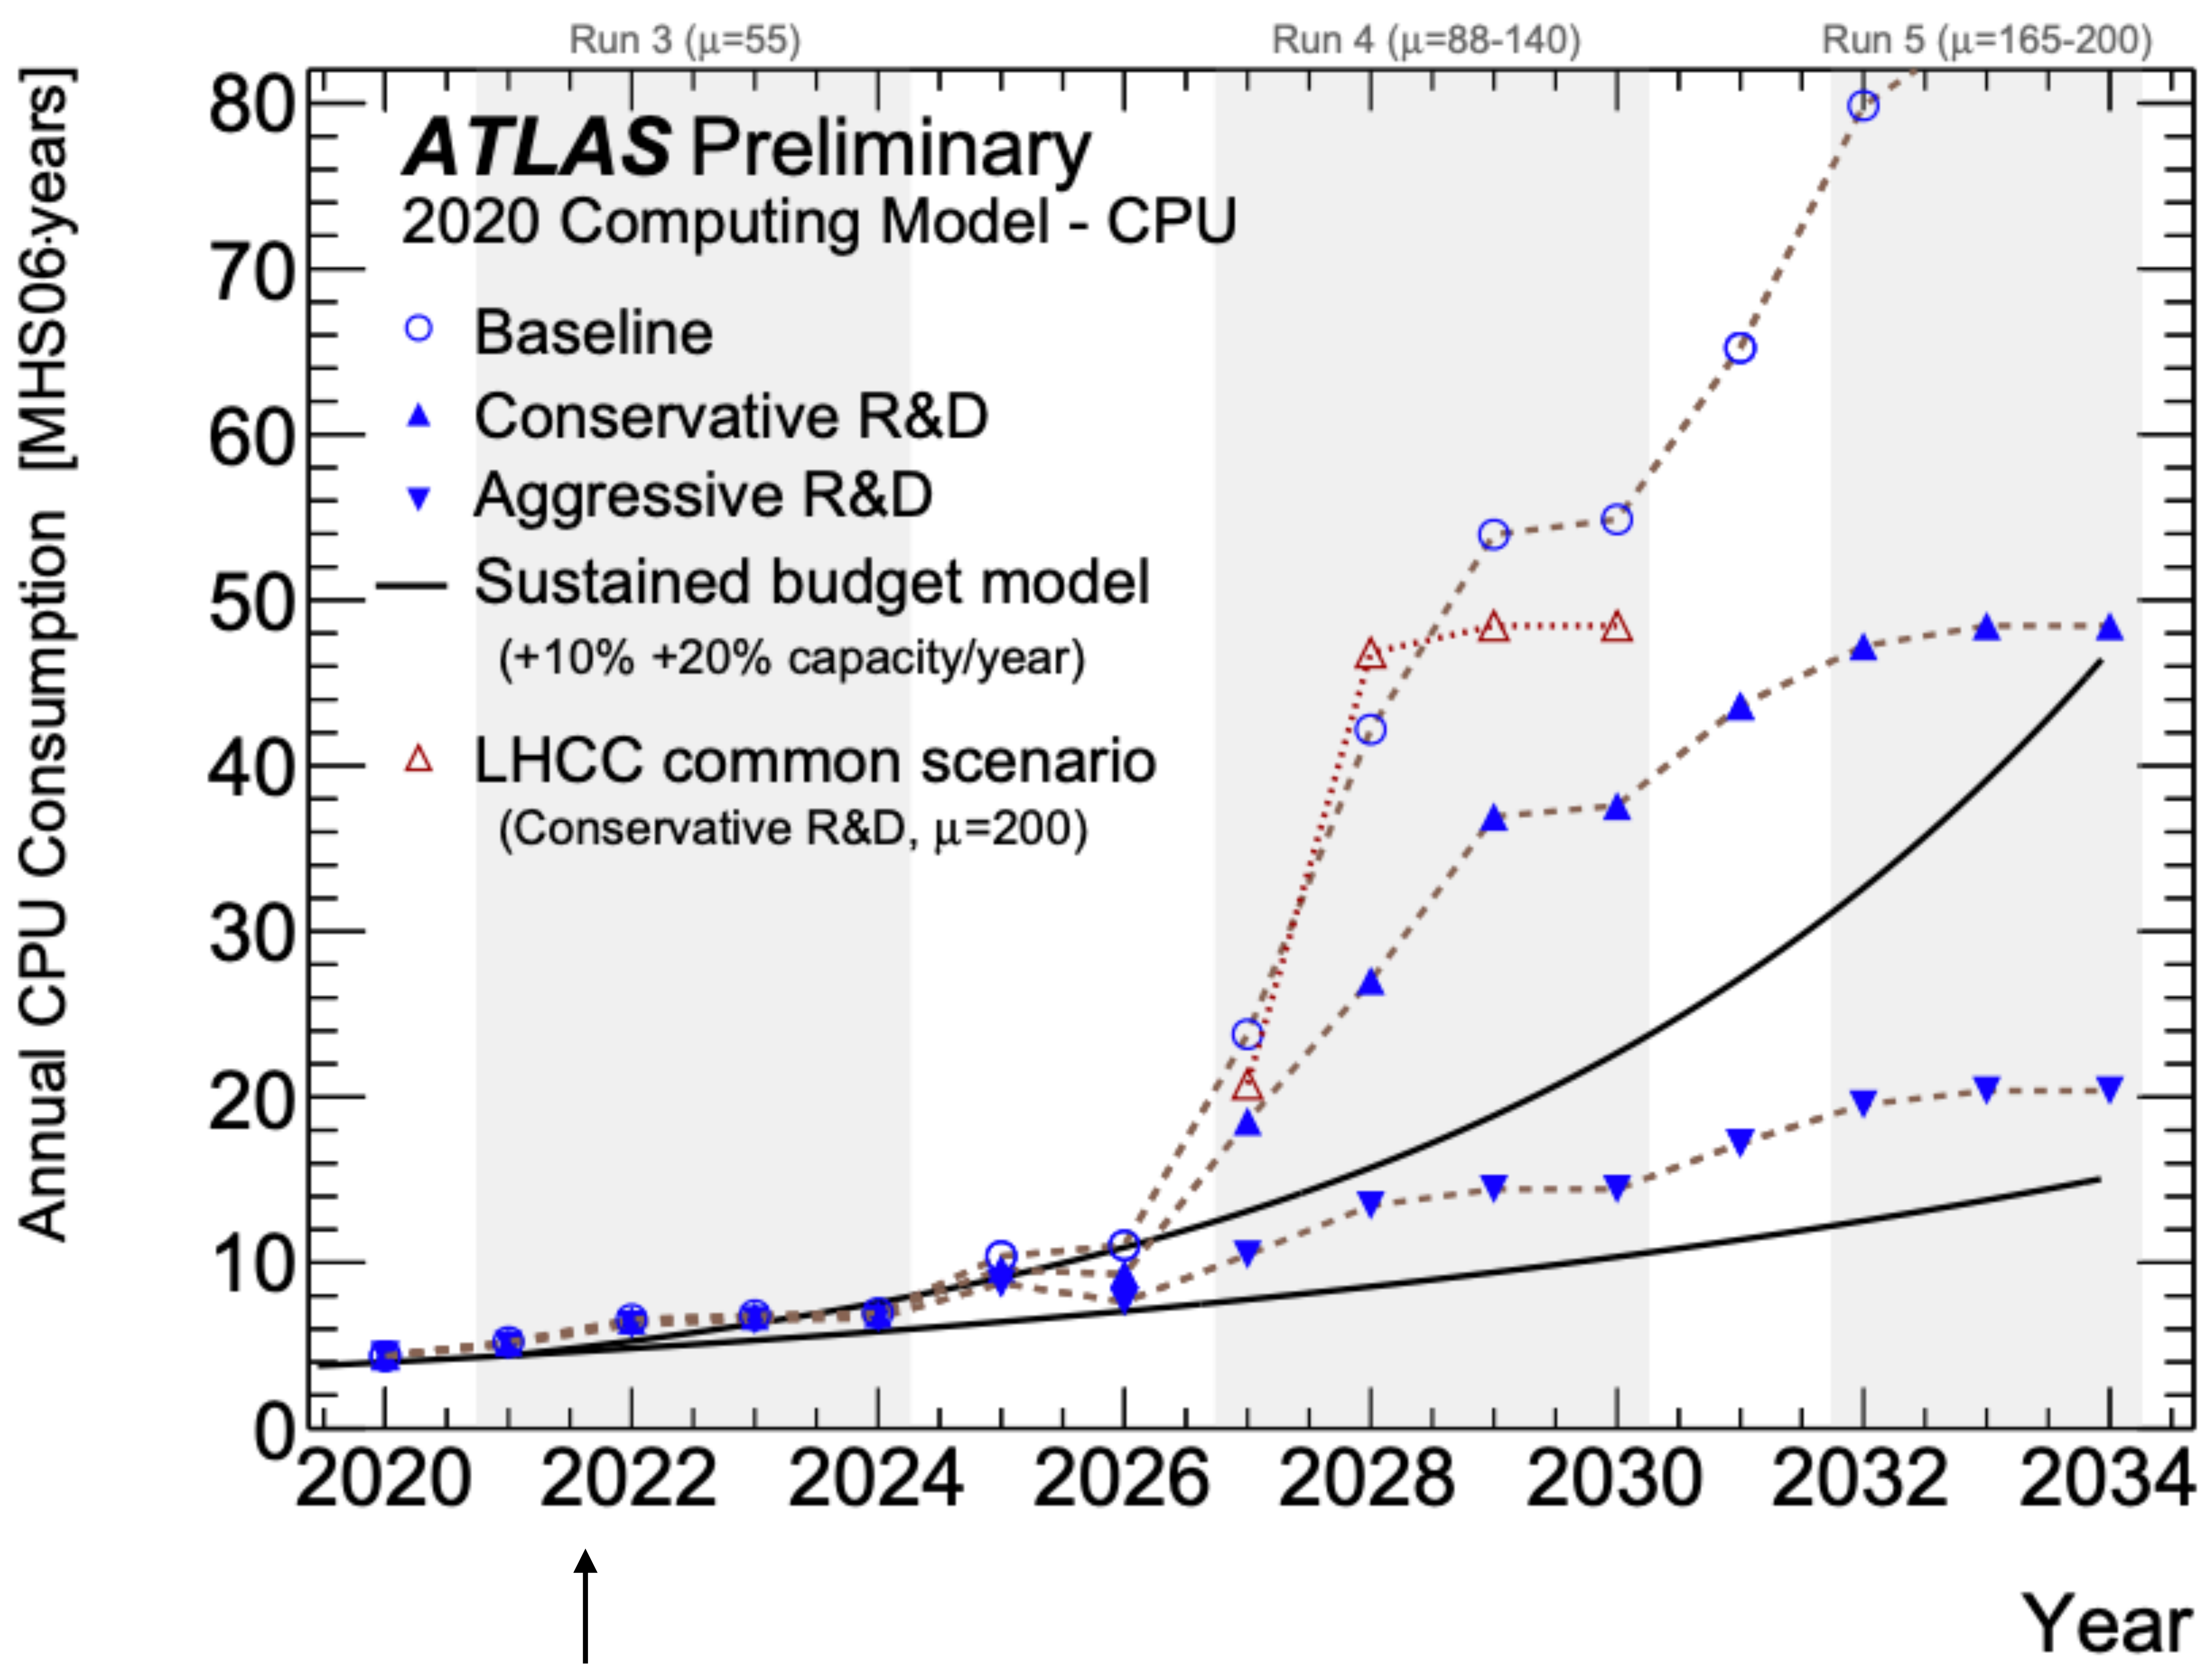

Just getting started

[ATLAS HL-LHC CDR]

We will collect 3x more data in Run-3, O(10x) more at the HL-LHC
What will Analysis Look like in that Future?

Hardware is changing

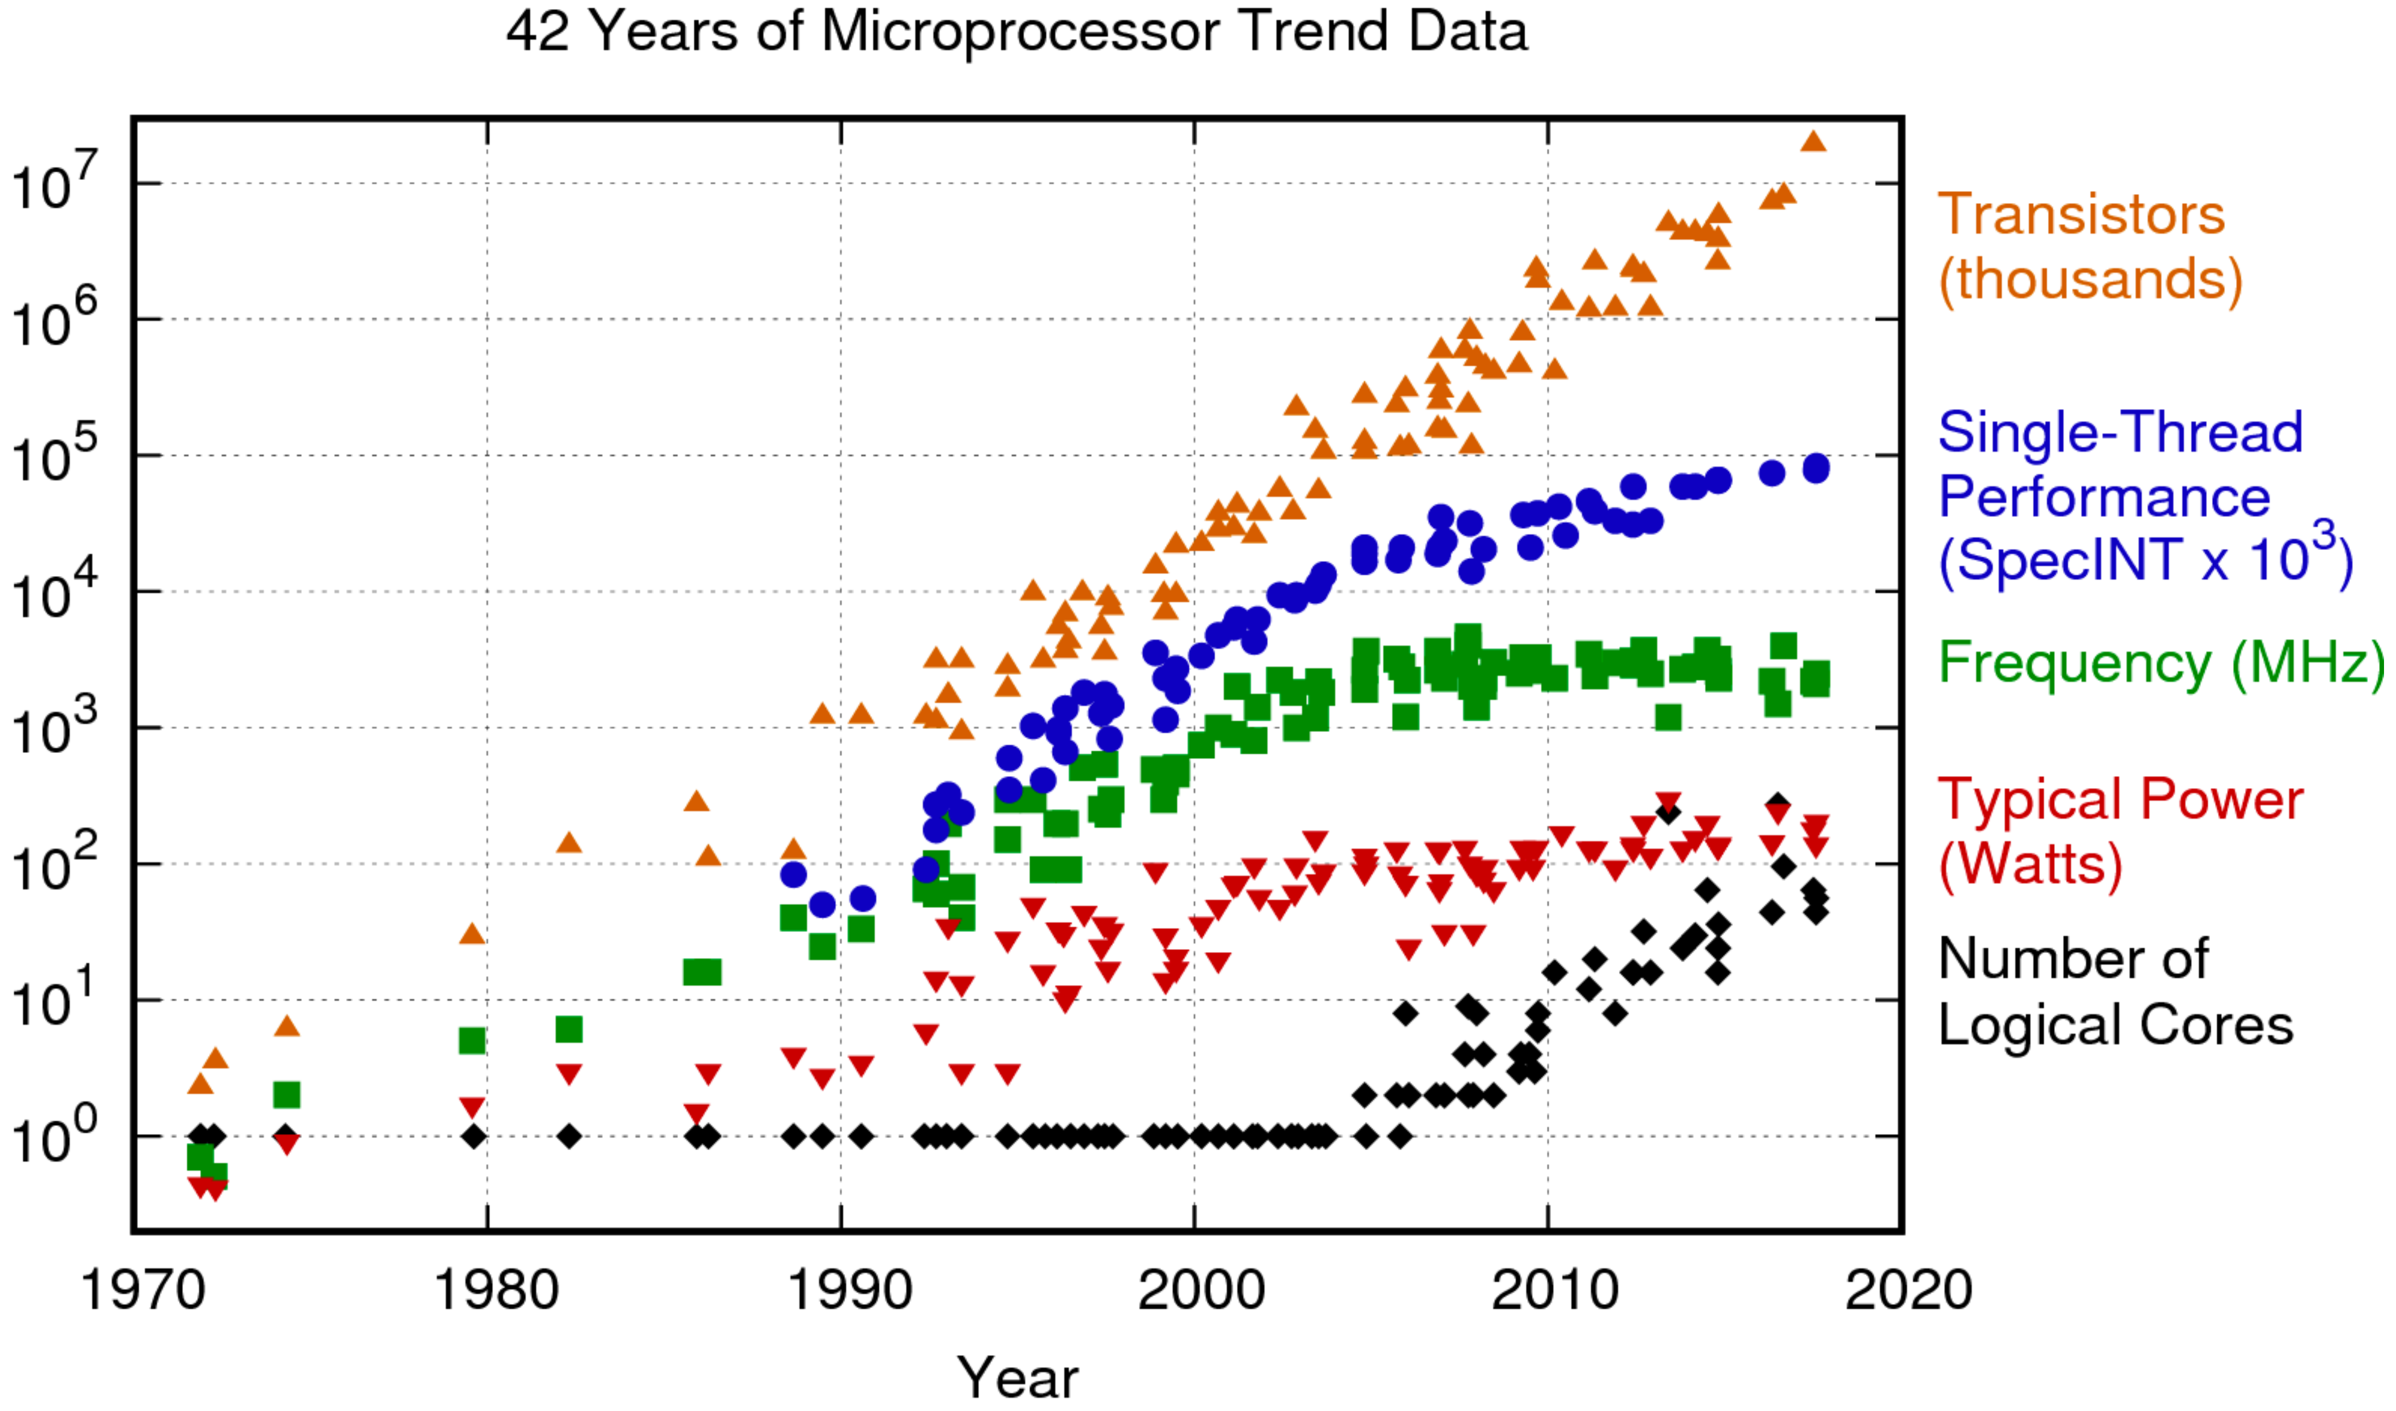

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

[K. Rupp]

turing lecture

DOI:10.1145/3282307

Innovations like domain-specific hardware, enhanced security, open instruction sets, and agile chip development will lead the way.

BY JOHN L. HENNESSY AND DAVID A. PATTERSON

A New Golden Age for Computer Architecture

WE BEGAN OUR Turing Lecture June 4, 2018[11] with a review of computer architecture since the 1960s. In addition to that review, here, we highlight current challenges and identify future opportunities, projecting another golden age for the field of computer architecture in the next decade, much like the 1980s when we did the research that led to our award, delivering gains in cost, energy, and security, as well as performance.

"Those who cannot remember the past are condemned to repeat it." —George Santayana, 1905

Software talks to hardware through a vocabulary called an instruction set architecture (ISA). By the early 1960s, IBM had four incompatible lines of computers, each with its own ISA, software stack, I/O system, and market niche—targeting small business, large business, scientific, and real time, respectively. IBM engineers, including ACM A.M. Turing Award laureate Fred Brooks, Jr., thought they could create a single ISA that would efficiently unify all four of these ISA bases.

They needed a technical solution for how computers as inexpensive as

key insights

- Software advances can inspire architecture innovation.
- Elevating the hardware/software interface creates opportunities for architecture innovation.
- The marketplace ultimately settles architecture debates.

48 COMMUNICATIONS OF THE ACM | FEBRUARY 2019 | VOL. 62 | NO. 2

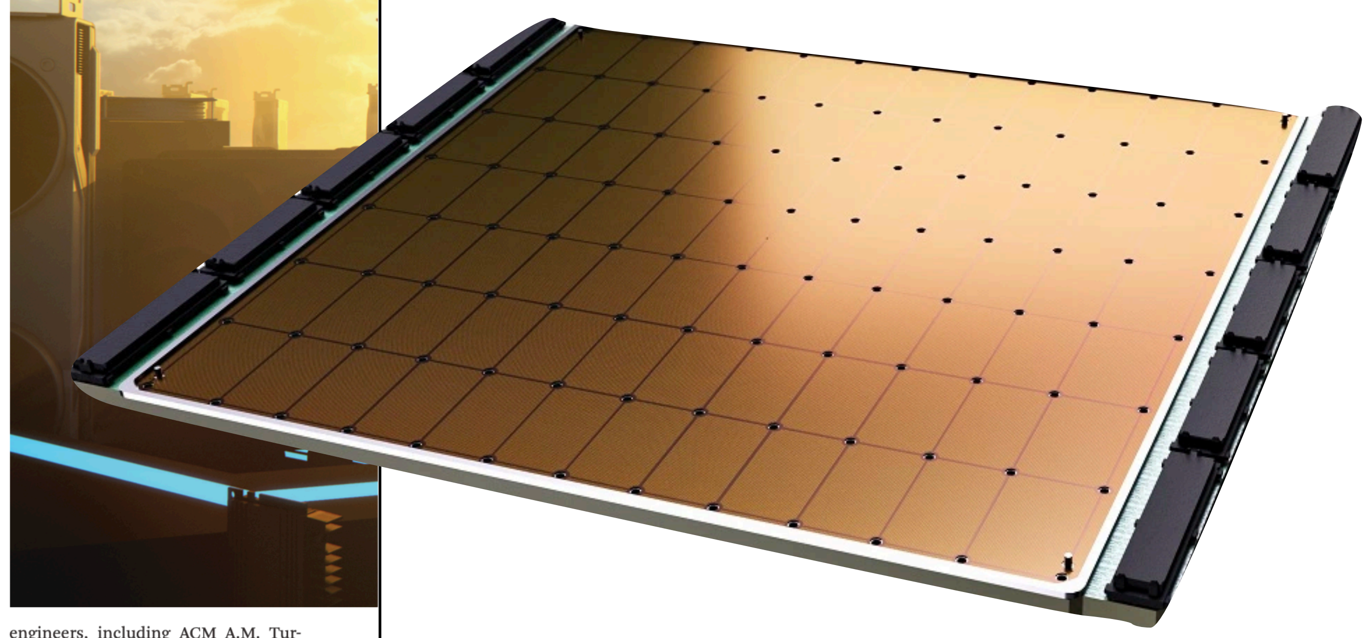

Waver-Scale Chips (Cerebras)

The age of general-purpose computer architectures is ending. Forced towards much more specialized chips. How will we deal with this?

Software is changing

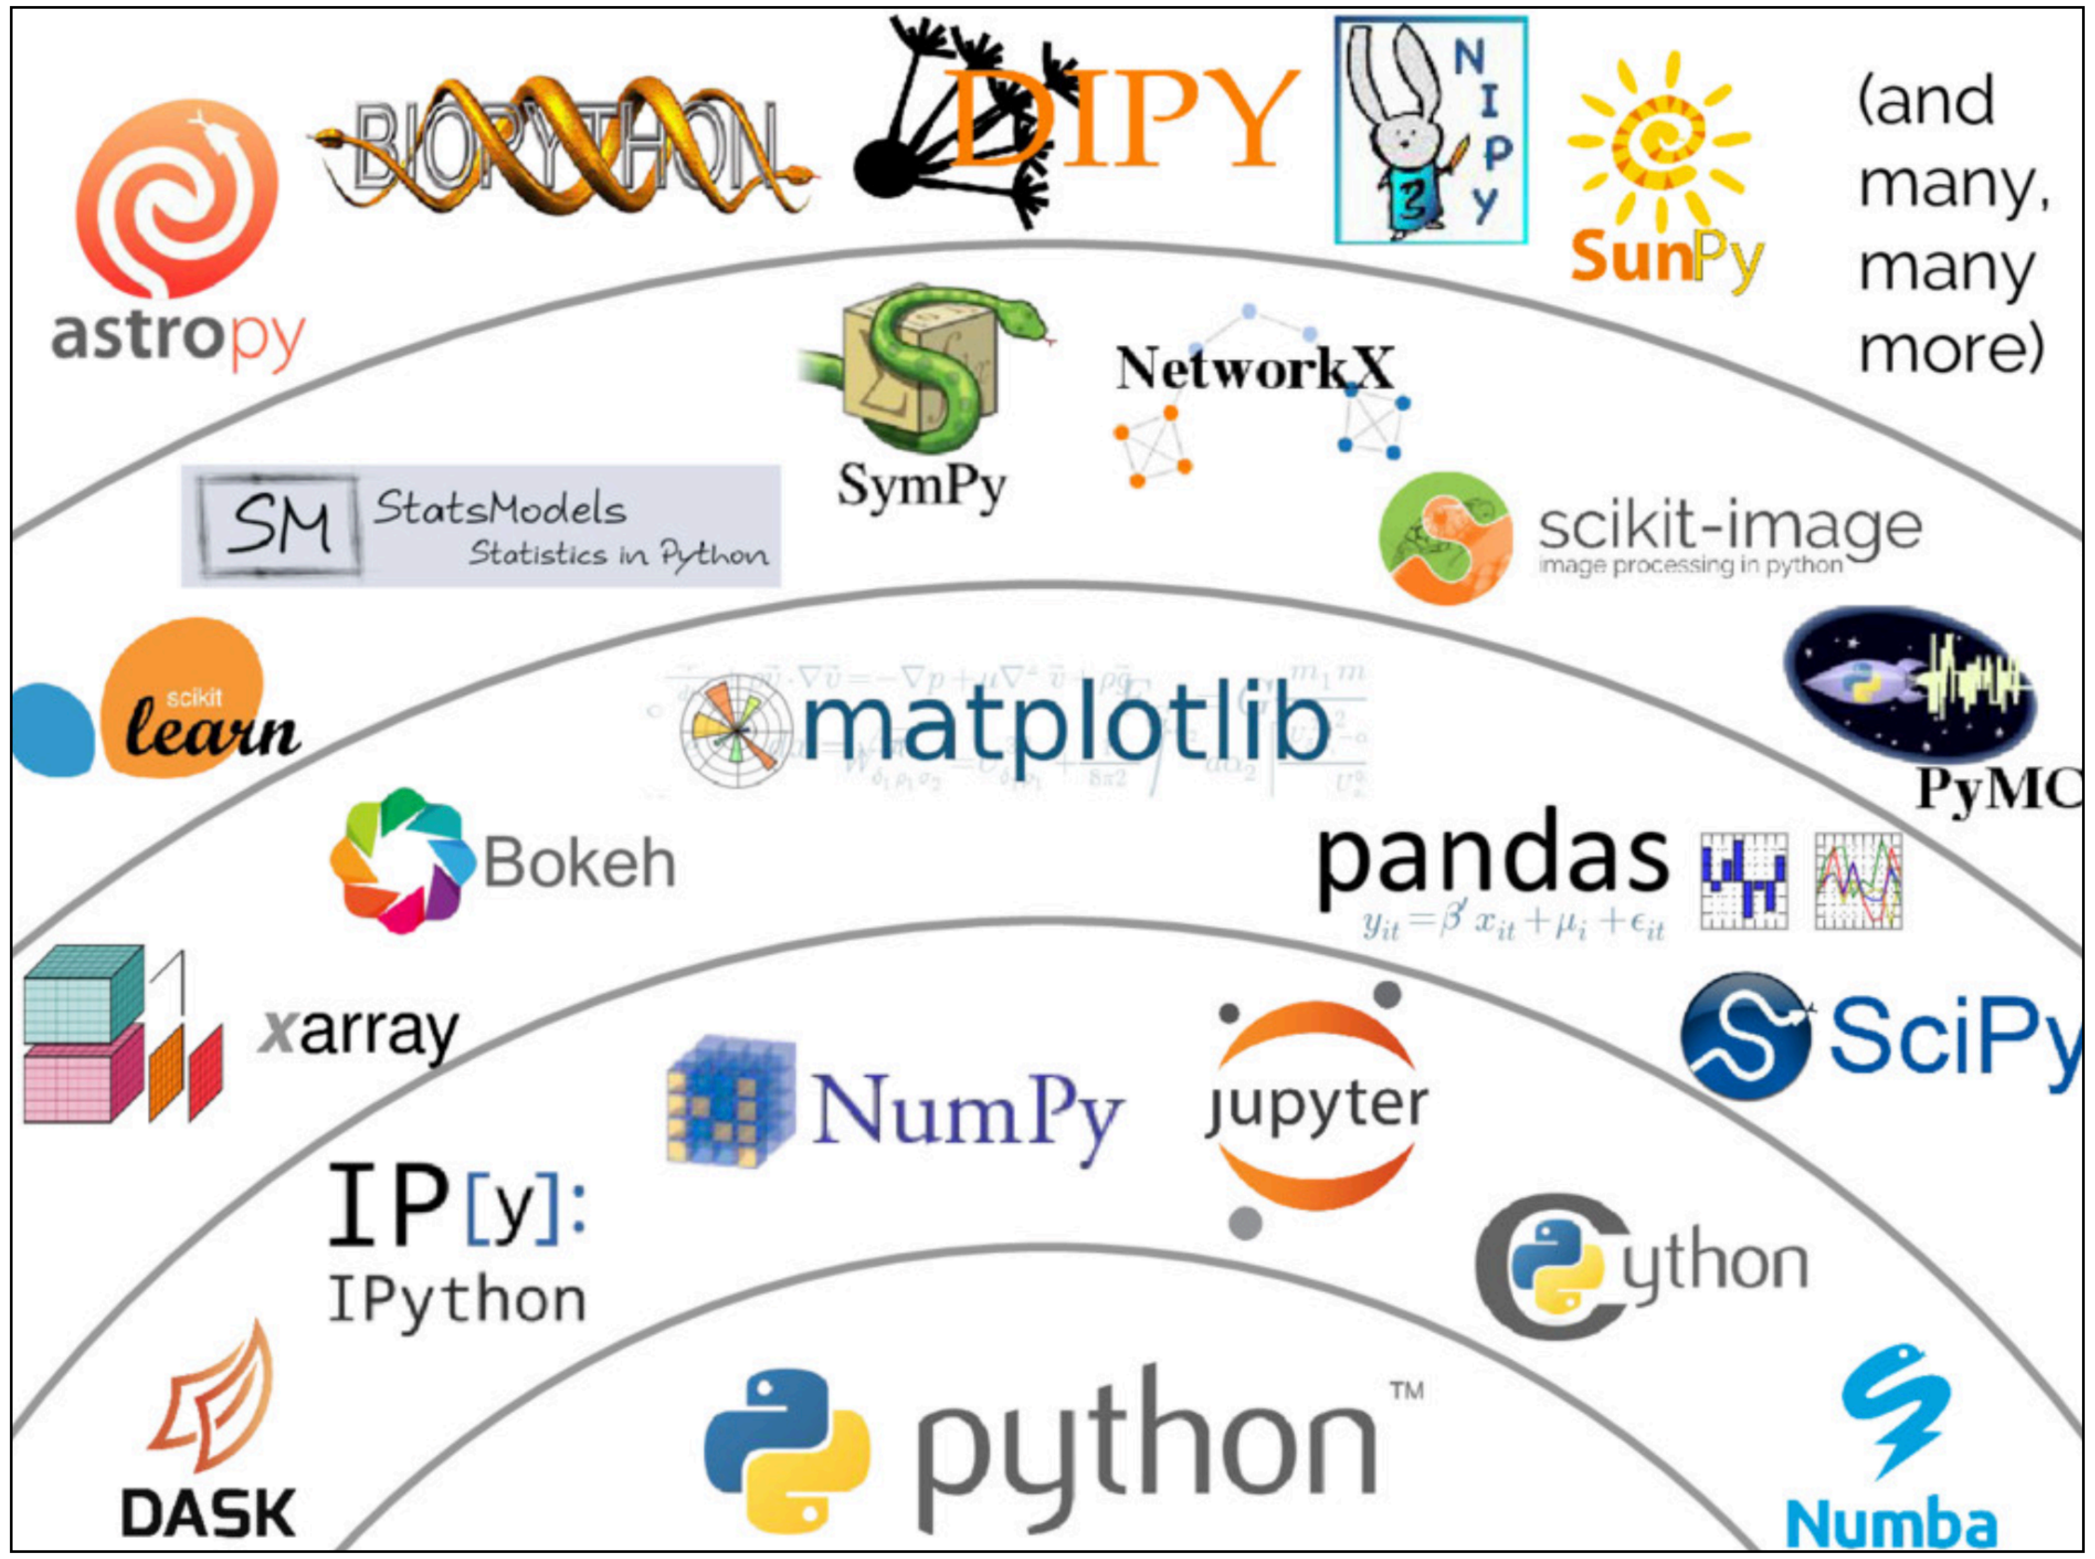

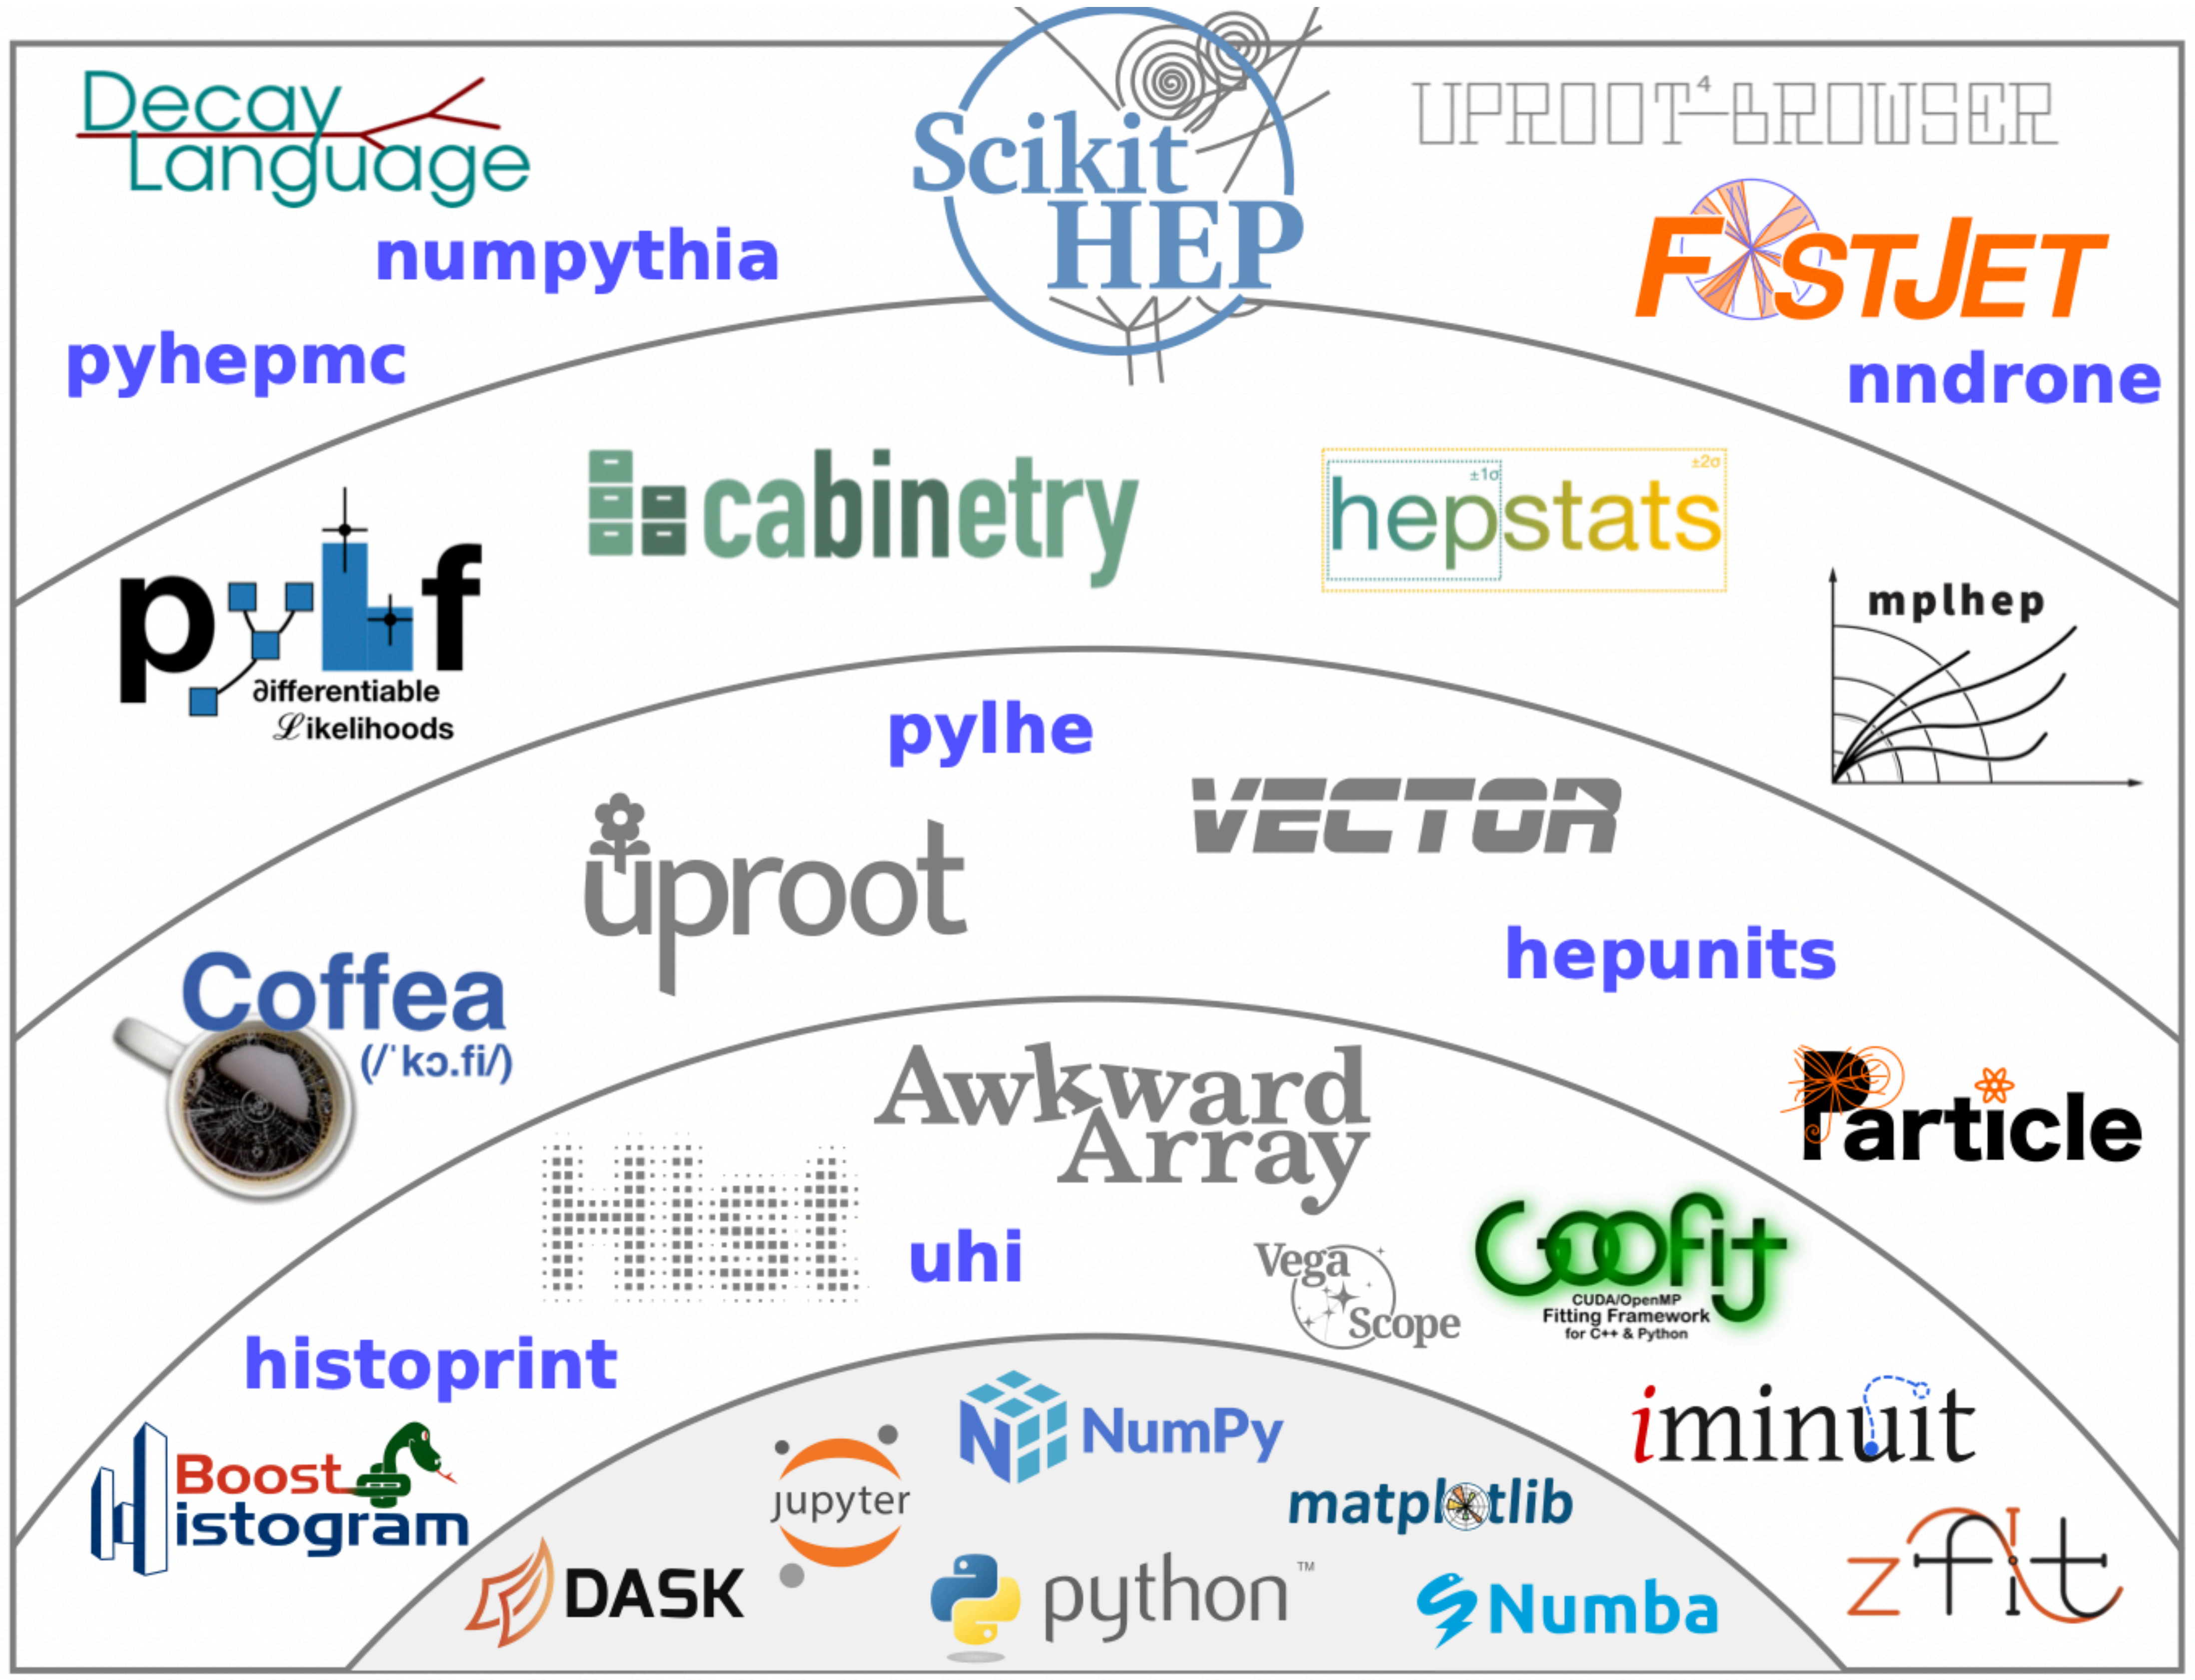

Data Science is much bigger than HNEP. Huge interest to use common tools for interoperability (the “I” in FAIR) & sustainability

Python currently dominant in that space.

Software is changing, maybe fundamentally

… the *other* breakthrough
event that happened in 2012

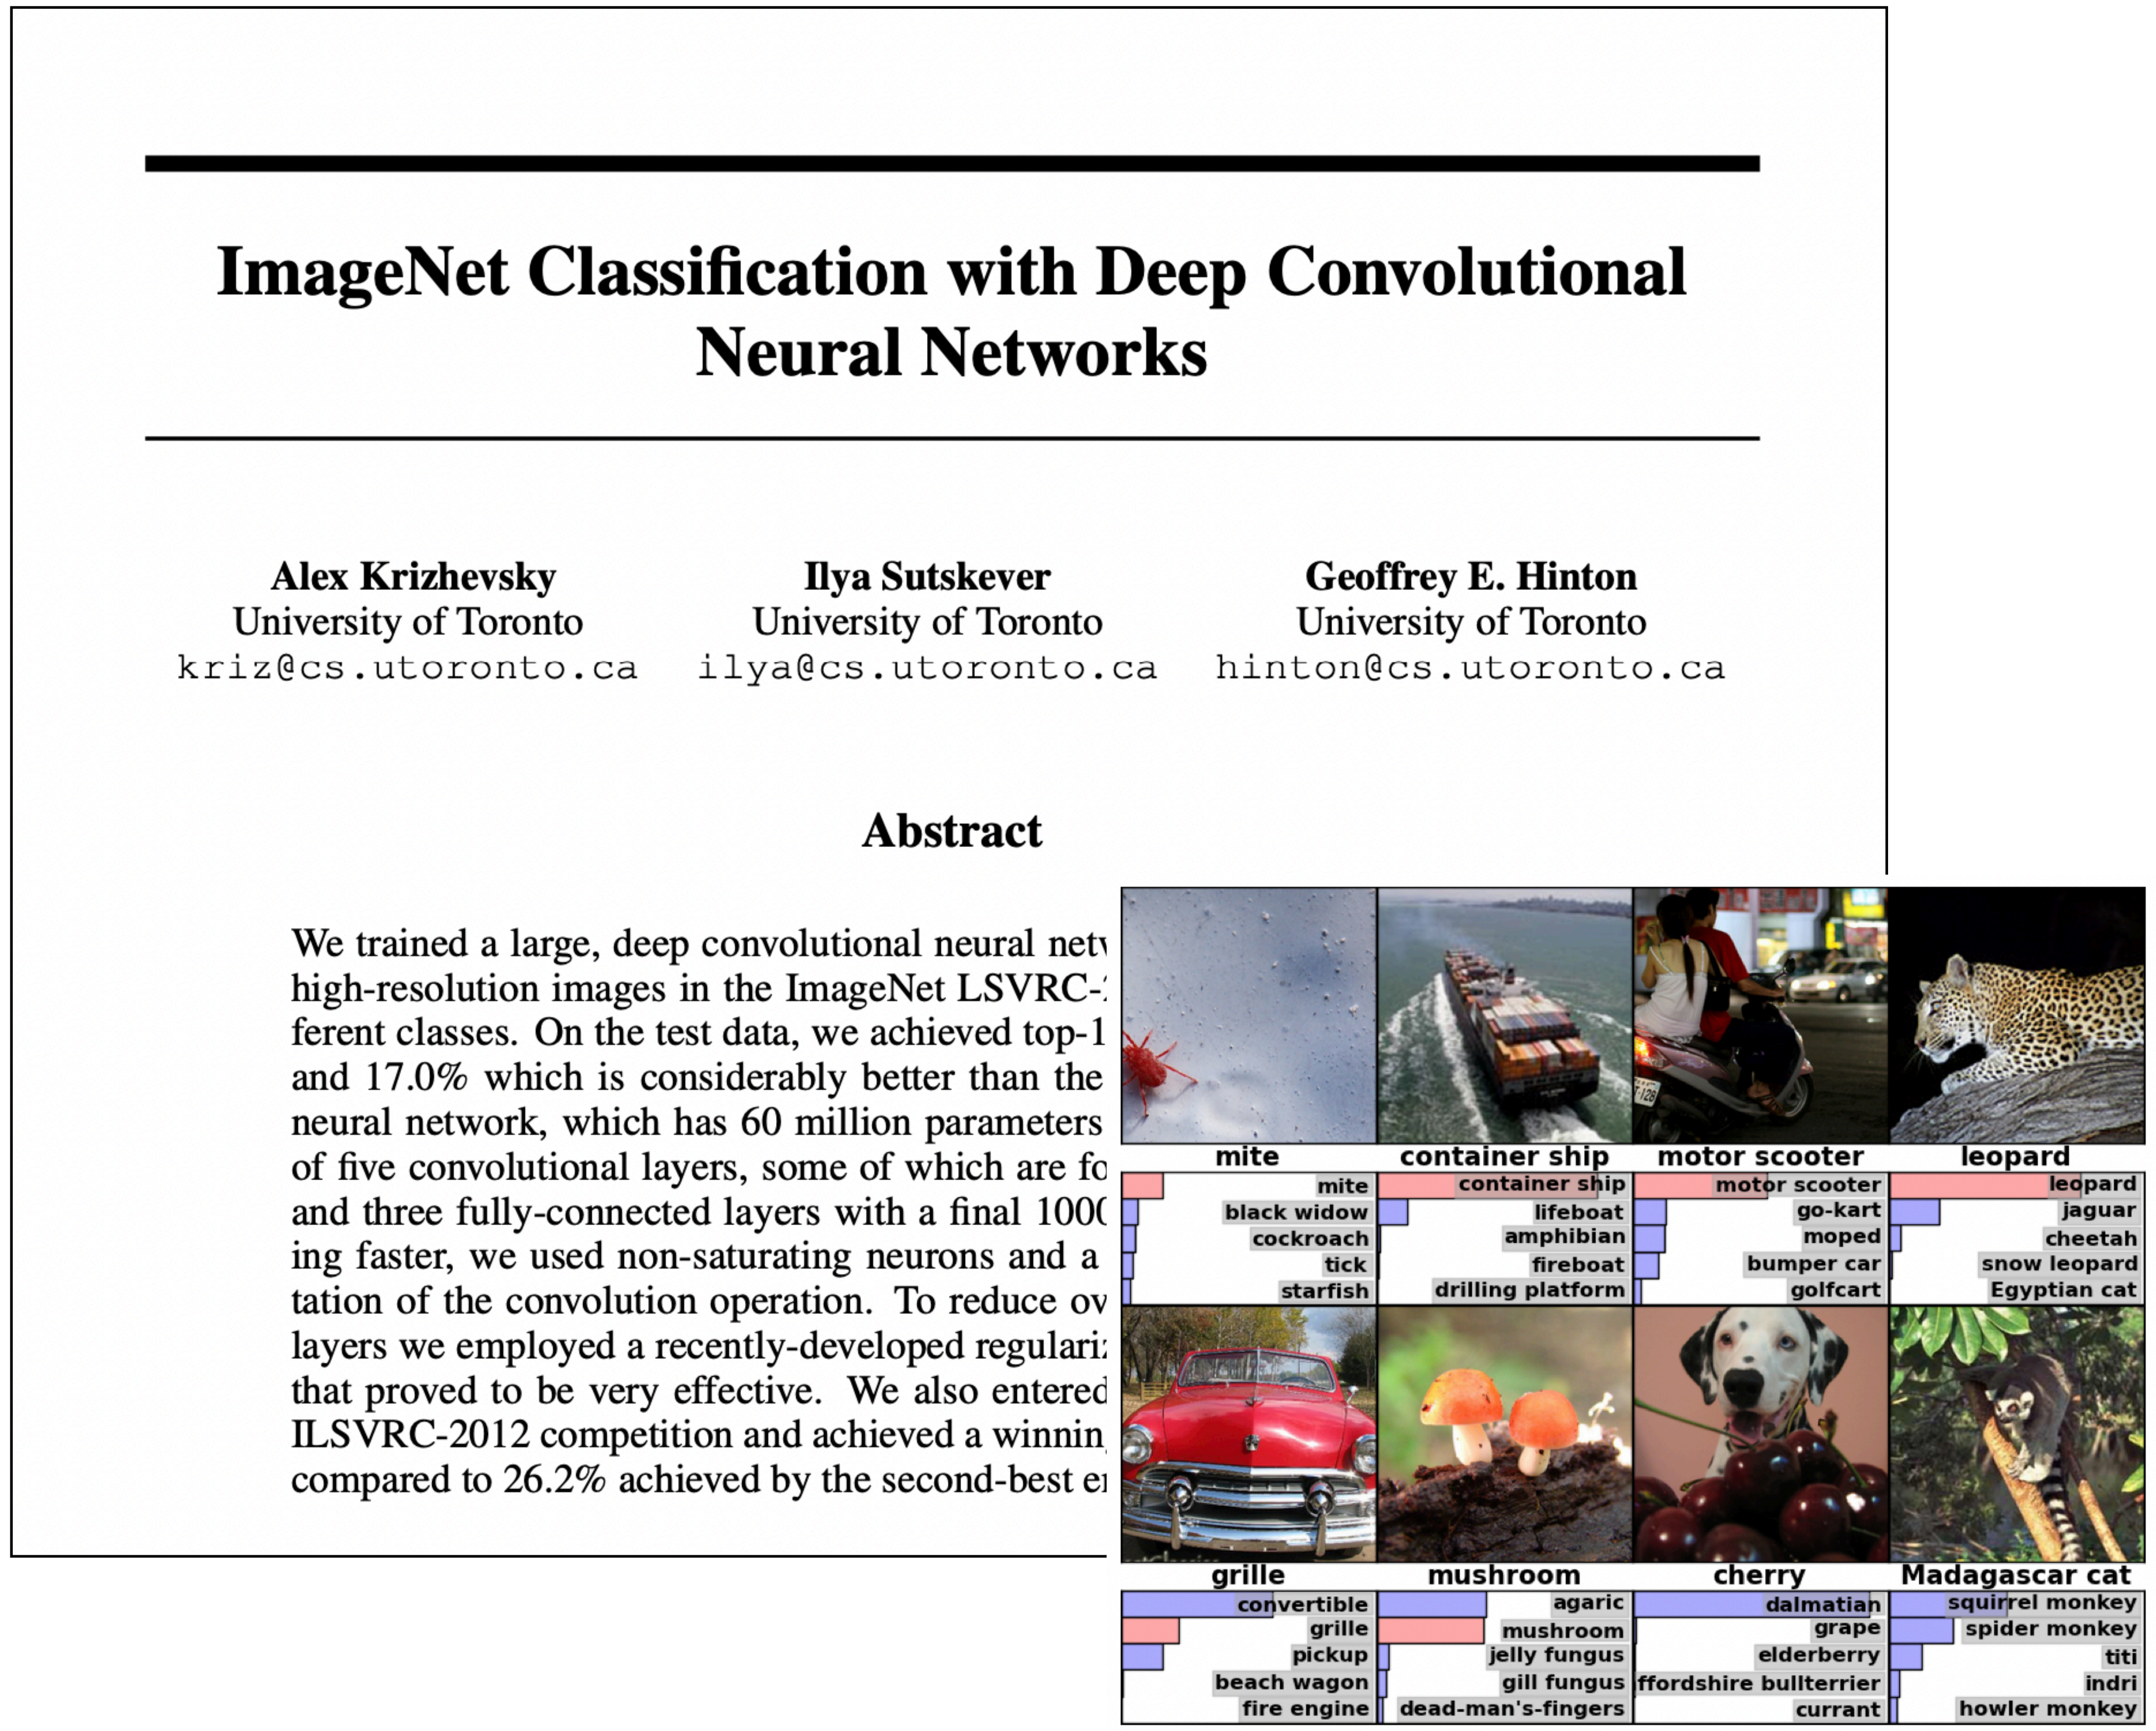

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural netv
high-resolution images in the ImageNet LSVRC-
ferent classes. On the test data, we achieved top-1
and 17.0% which is considerably better than the
neural network, which has 60 million parameters
of five convolutional layers, some of which are fo
and three fully-connected layers with a final 100(
ing faster, we used non-saturating neurons and a
tation of the convolution operation. To reduce ov
layers we employed a recently-developed regulari
that proved to be very effective. We also entered
ILSVRC-2012 competition and achieved a winnin
compared to 26.2% achieved by the second-best e

Deep Learning takes over ML
thanks to big datasets and GPUs

ML as a **new way to write software** by providing data rather than code
What is the "software development" workflow w/o source code?

Infrastructure is changing

We're not alone in global, federated computation anymore

Google Cloud

aws

Alongside by very successful open source software projects

openstack

docker

kubernetes

Can we use this in a way that's favourable to us?

Analysis Facilities, next-Gen Grid Sites?

In light of these, it's clear that infrastructure in HEP & other fundamental sciences will shift. A possible wish-list:

- is this deployable at scale? reproducibility at infrastructure level?

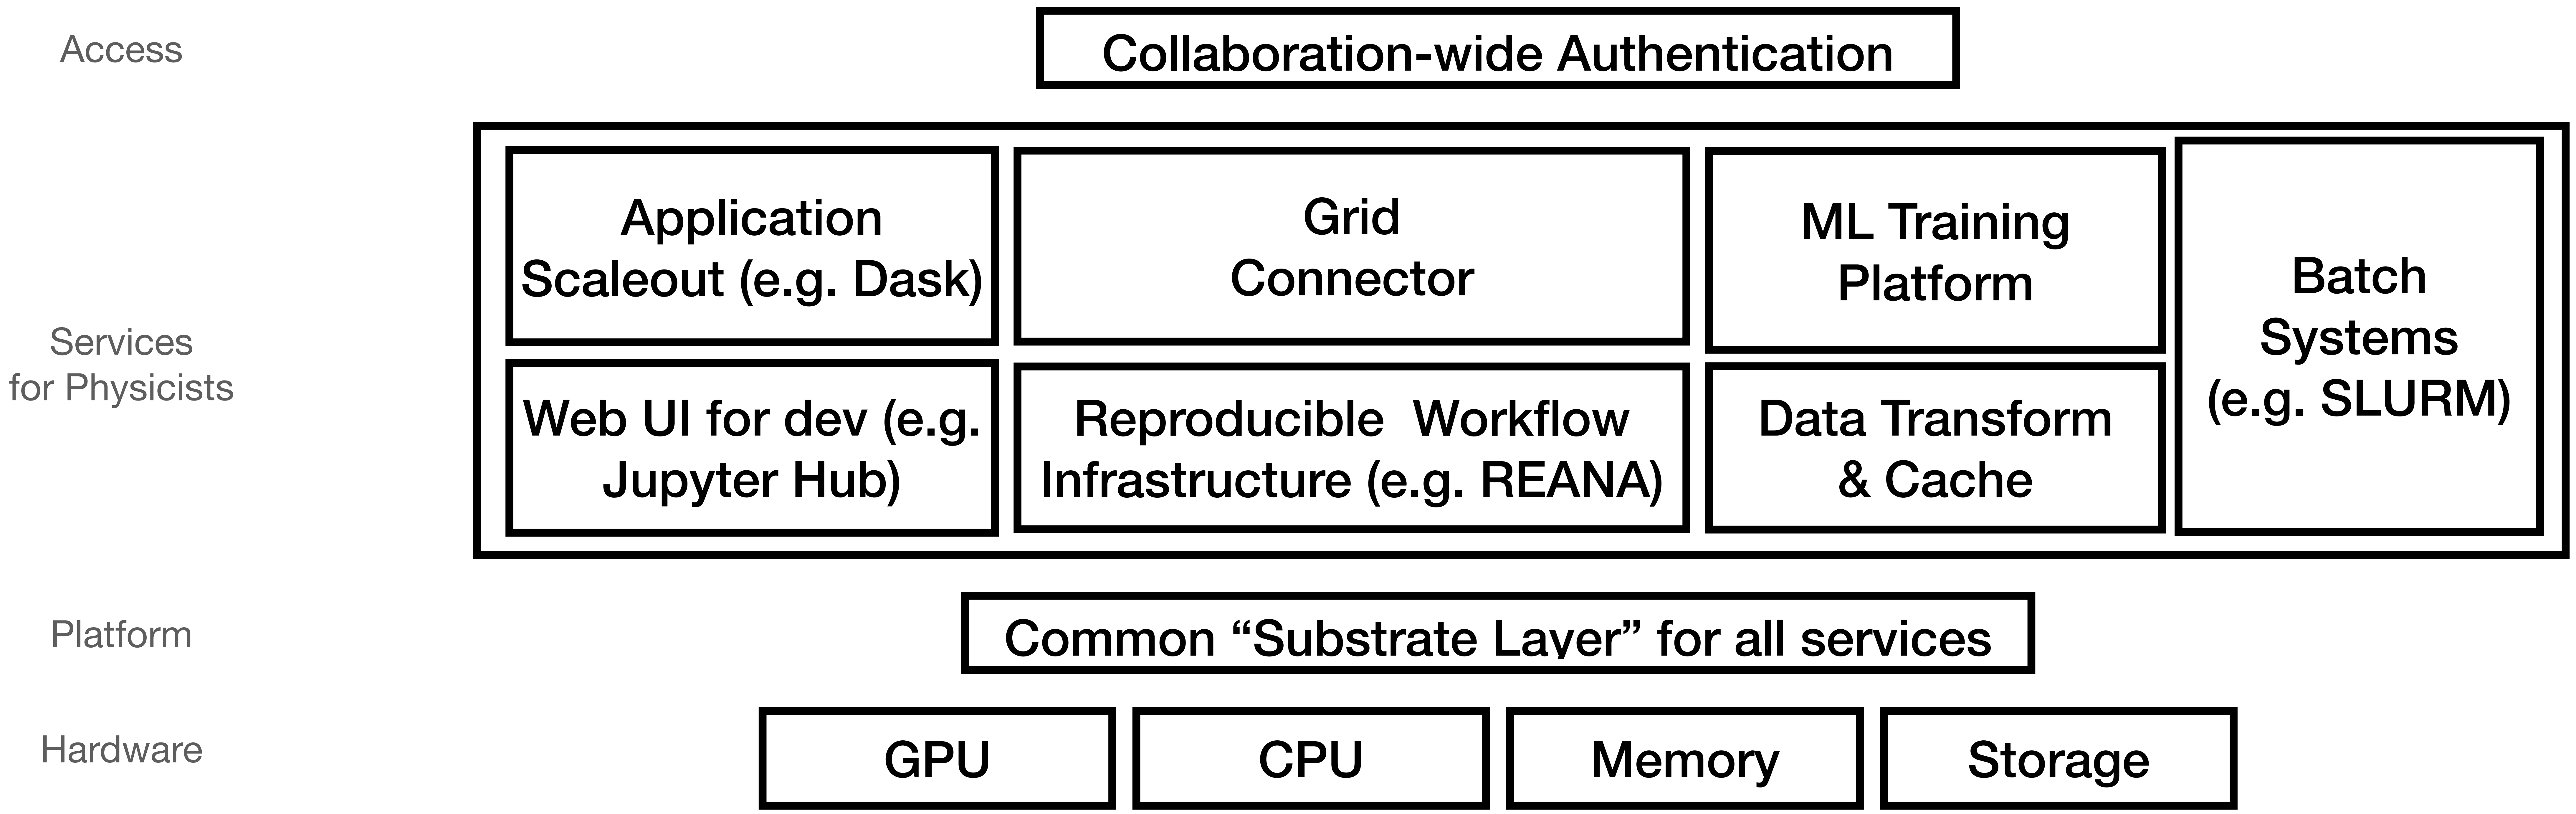

ATLAS-Google Cloud Project

Cloud Advantage: can test before committing large in-house resources

- collaboration between ATLAS & Google to explore what new infrastructure could look like

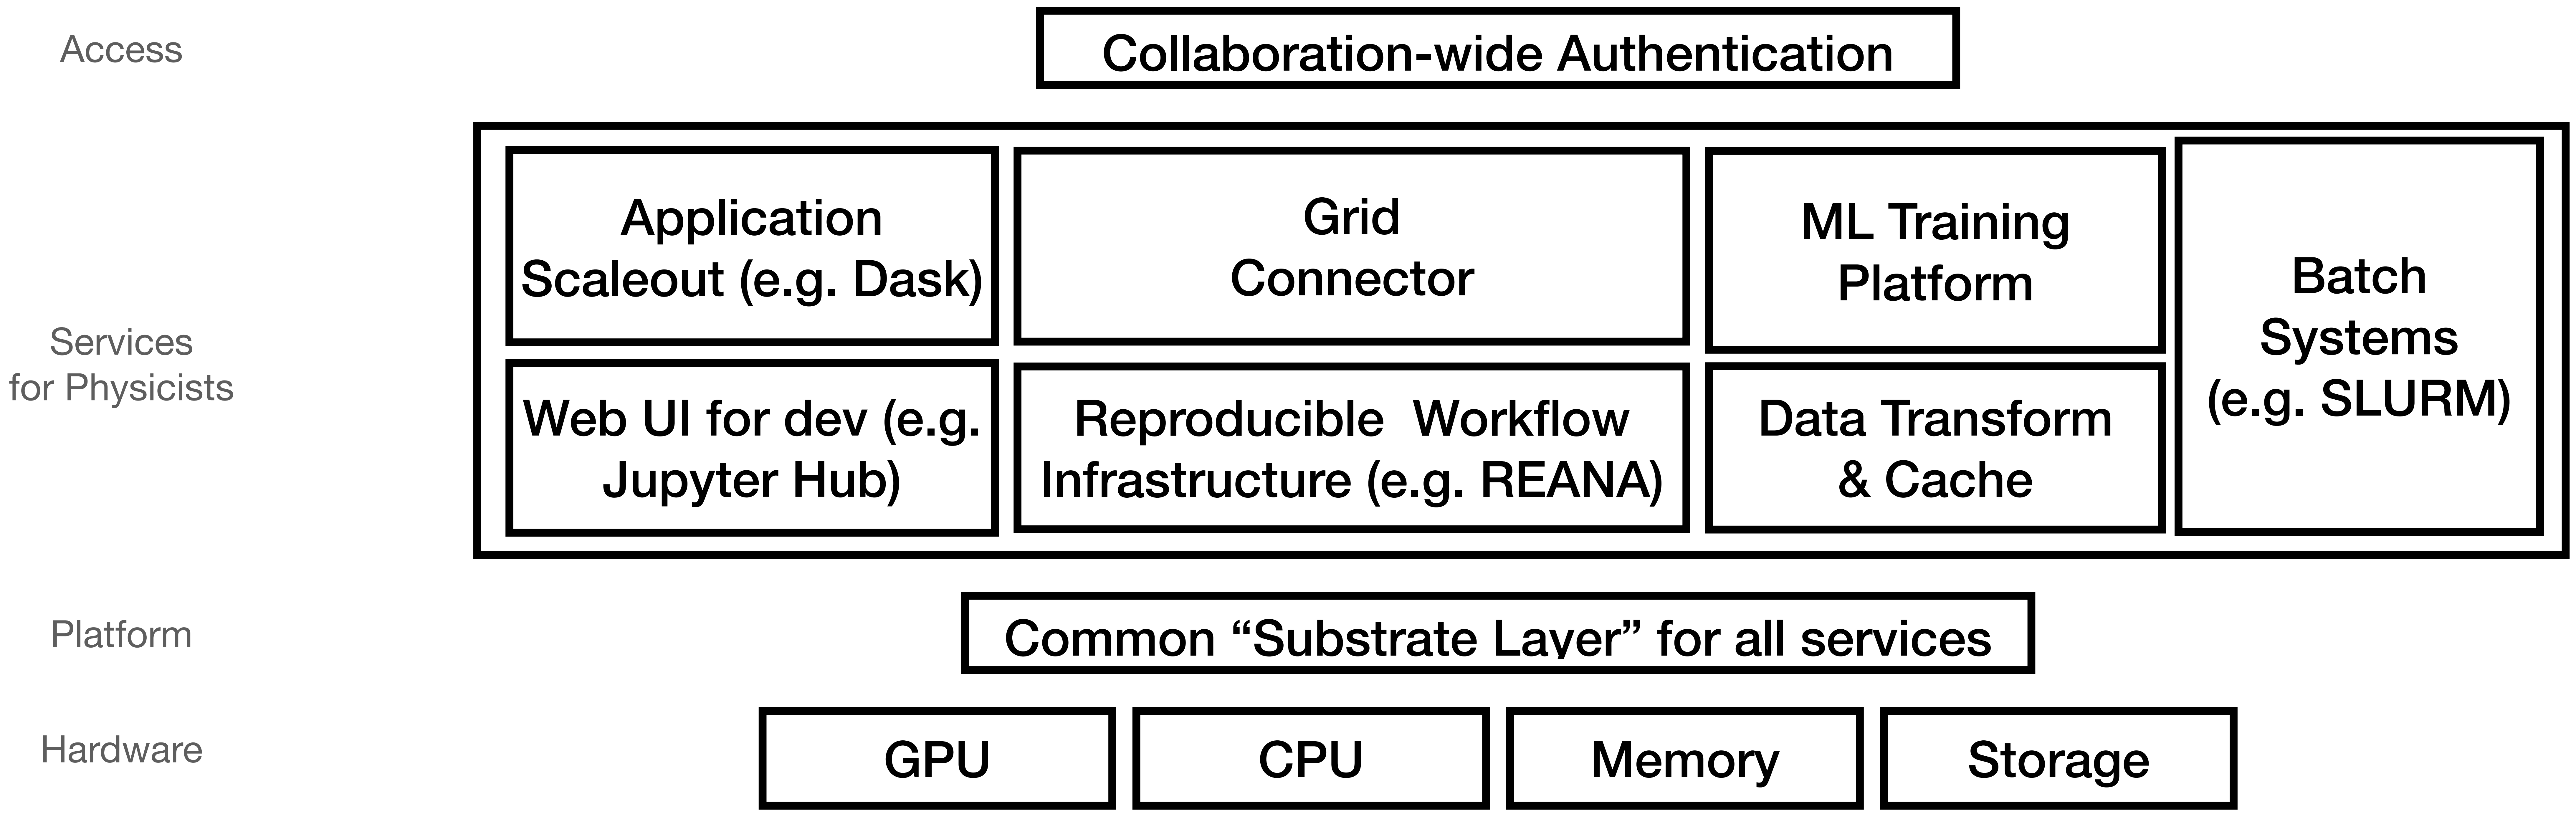

Cloud Characteristics

Important to keep in mind key differences to on-premise computing

- cloud is <u>very</u> elastic: 50k cores for 1h == 1k cores for 50h
- esp. useful for applications where wall clock time important (e.g. interactive analysis, seqeuential algorithms, …)

Cloud Characteristics

Important to keep in mind key differences to on-premise computing

- cloud is (can be) very expensive if you don't watch out
- use elasticity to scale on-demand, minimize data transfer
- avoid proprietary offerings to avoid lock in - stay open source

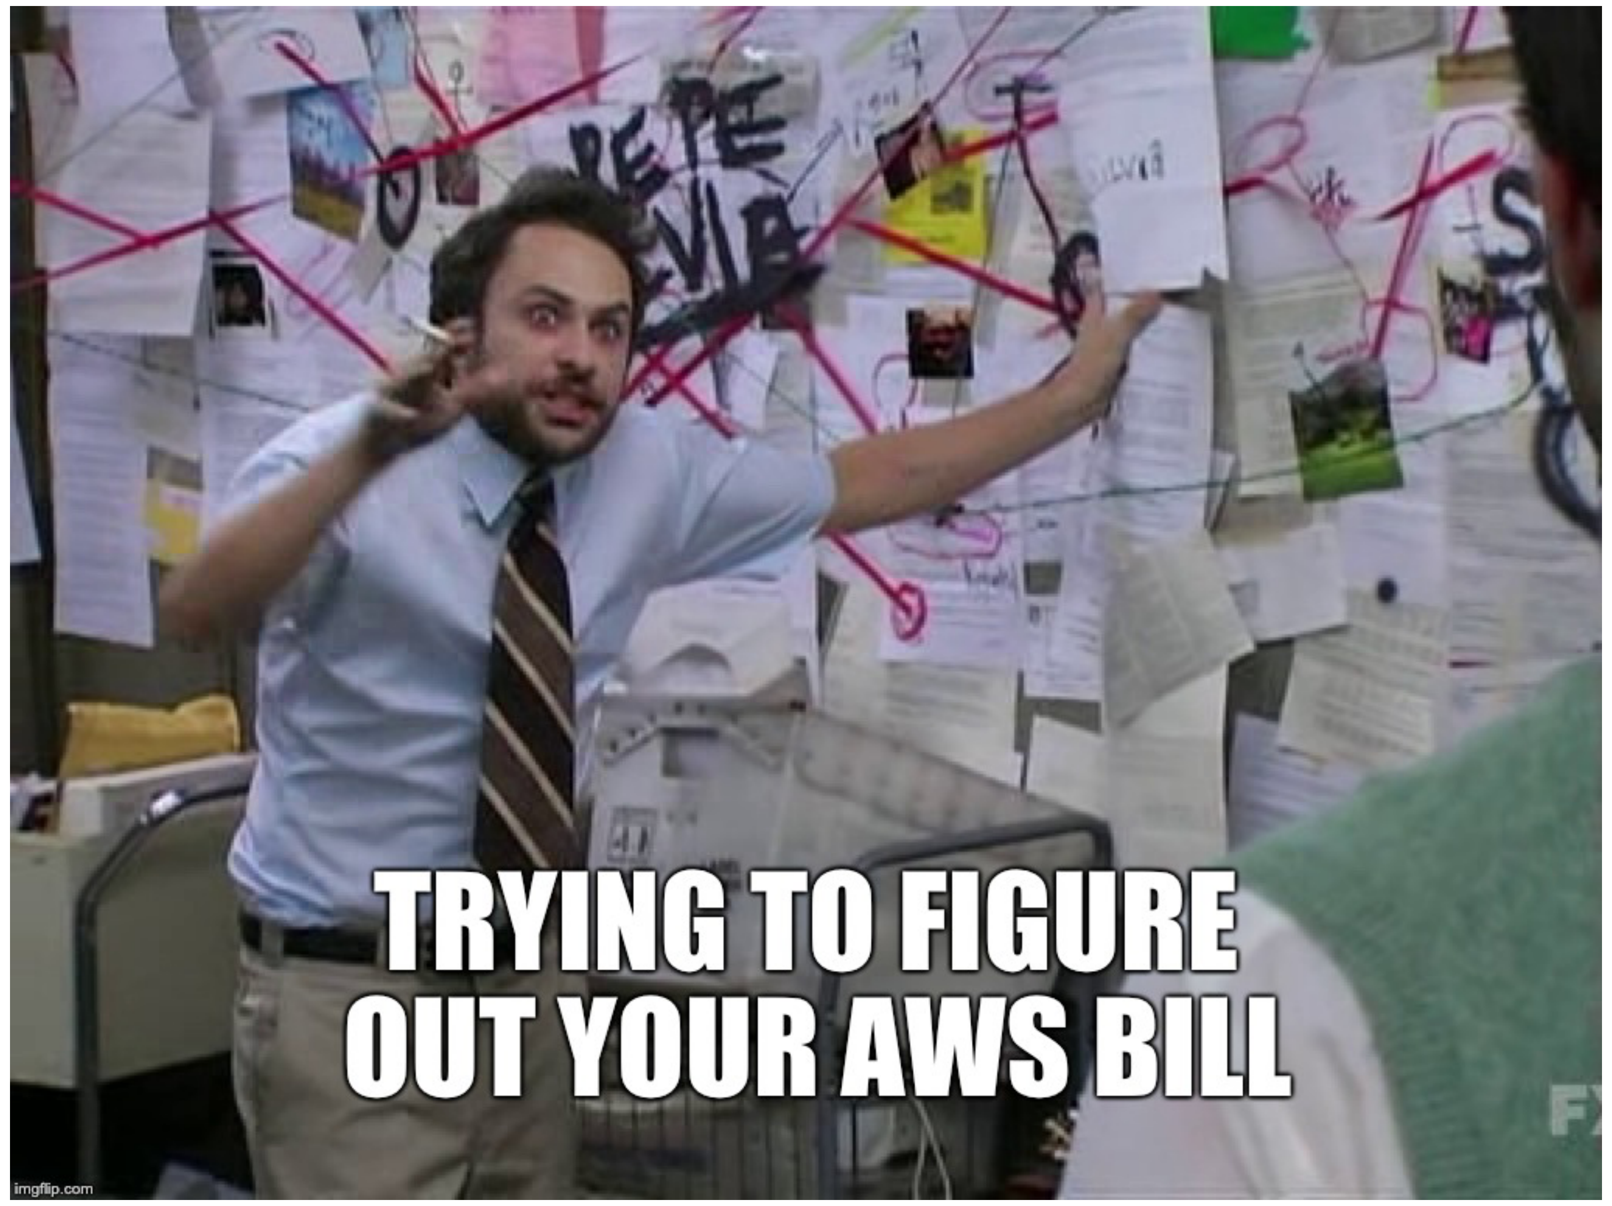

Taking Steps toward the Vision

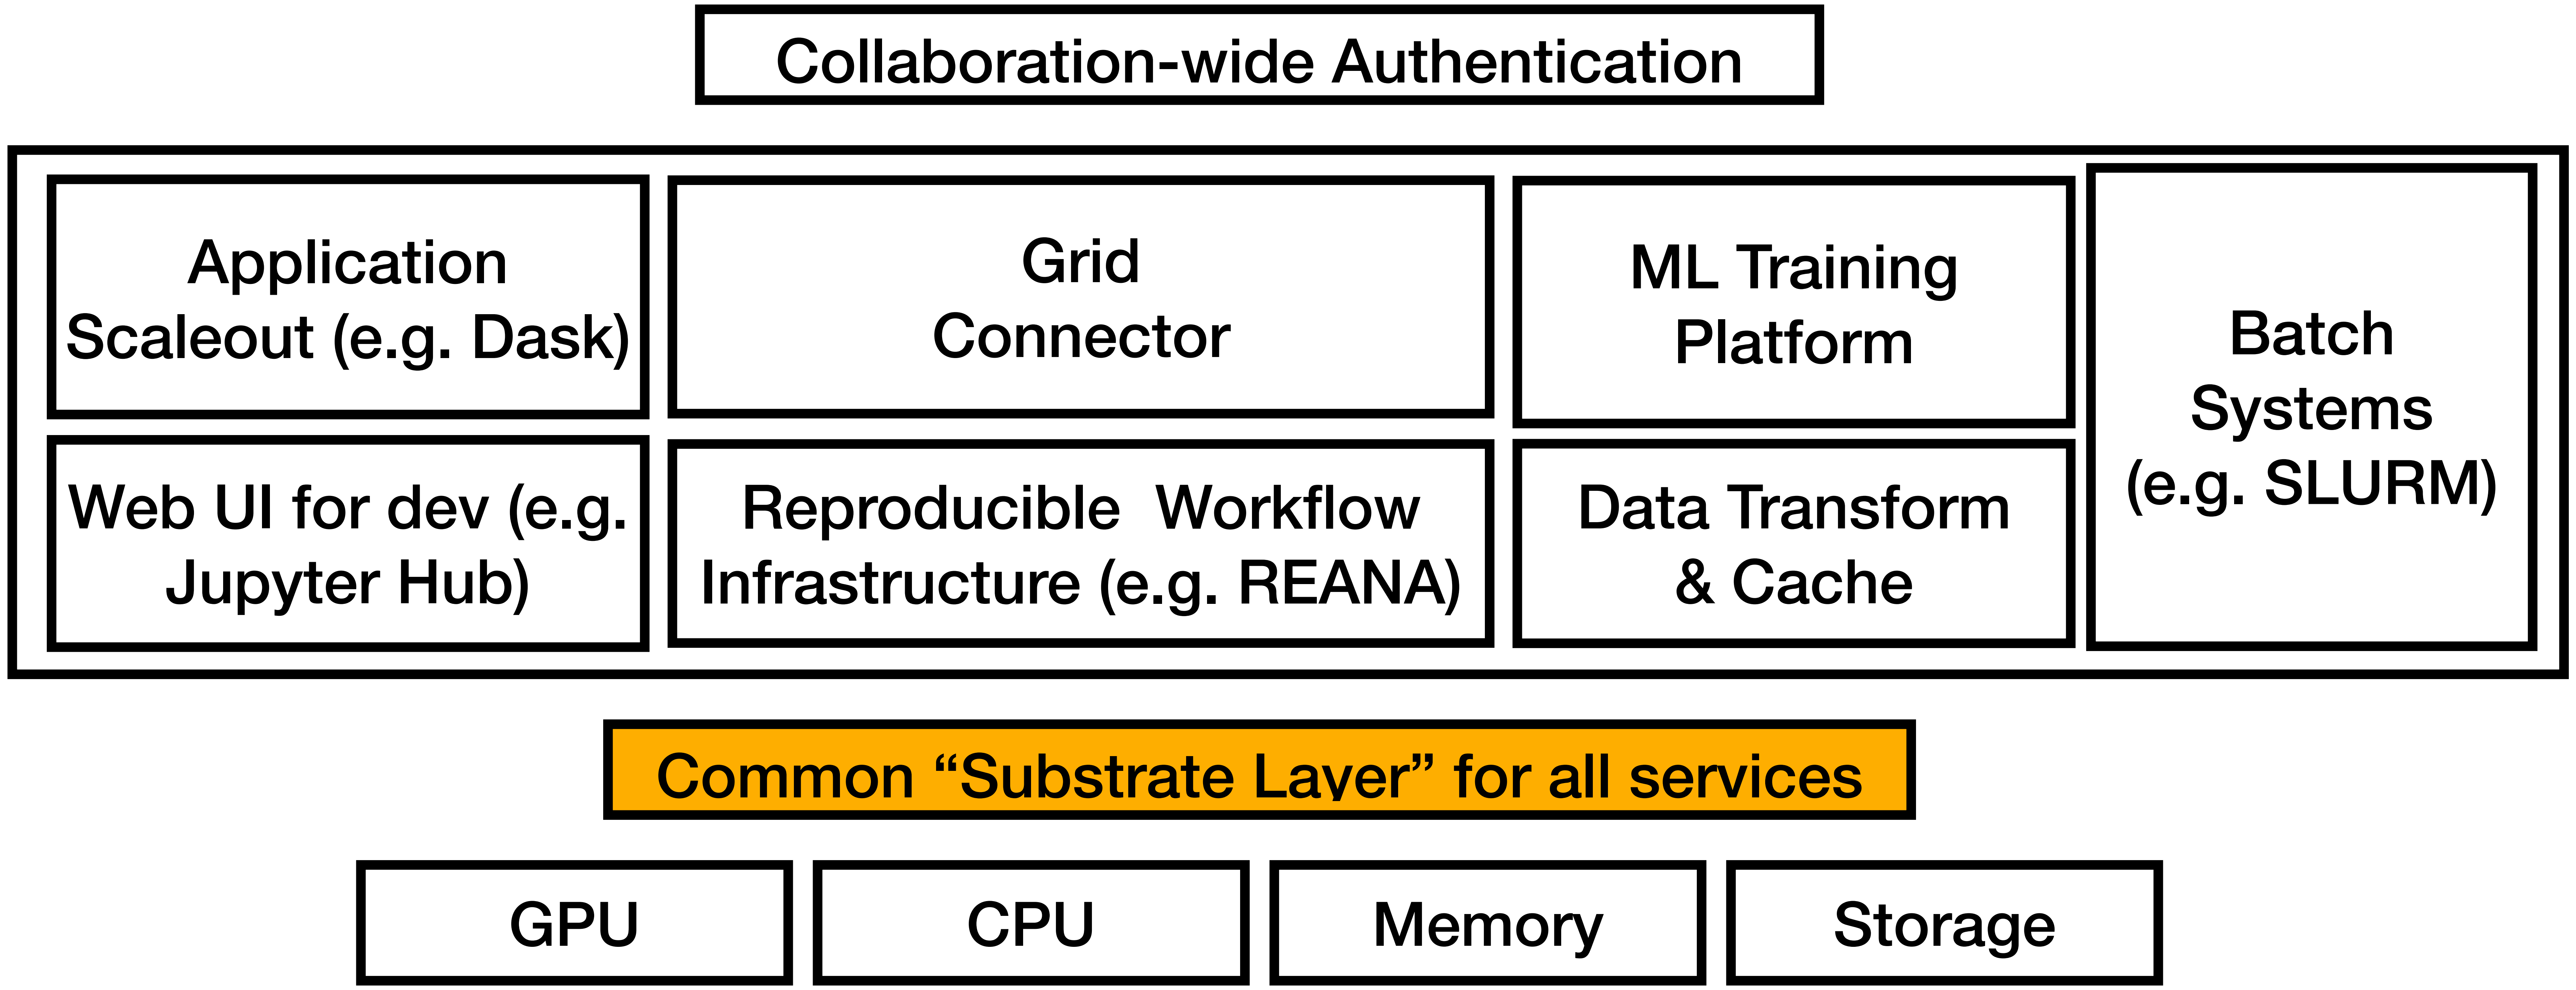

Platform Choice

Kubernetes is likely the best choice for new infrastructure projects.

- "Lingua Franca" of the cloud, transferable to our own infrastructure
- Target platform for many new projects starting today
- Similar to how we settled on Linux for the original Grid?

Universities / Labs

Cloud A

Cloud B

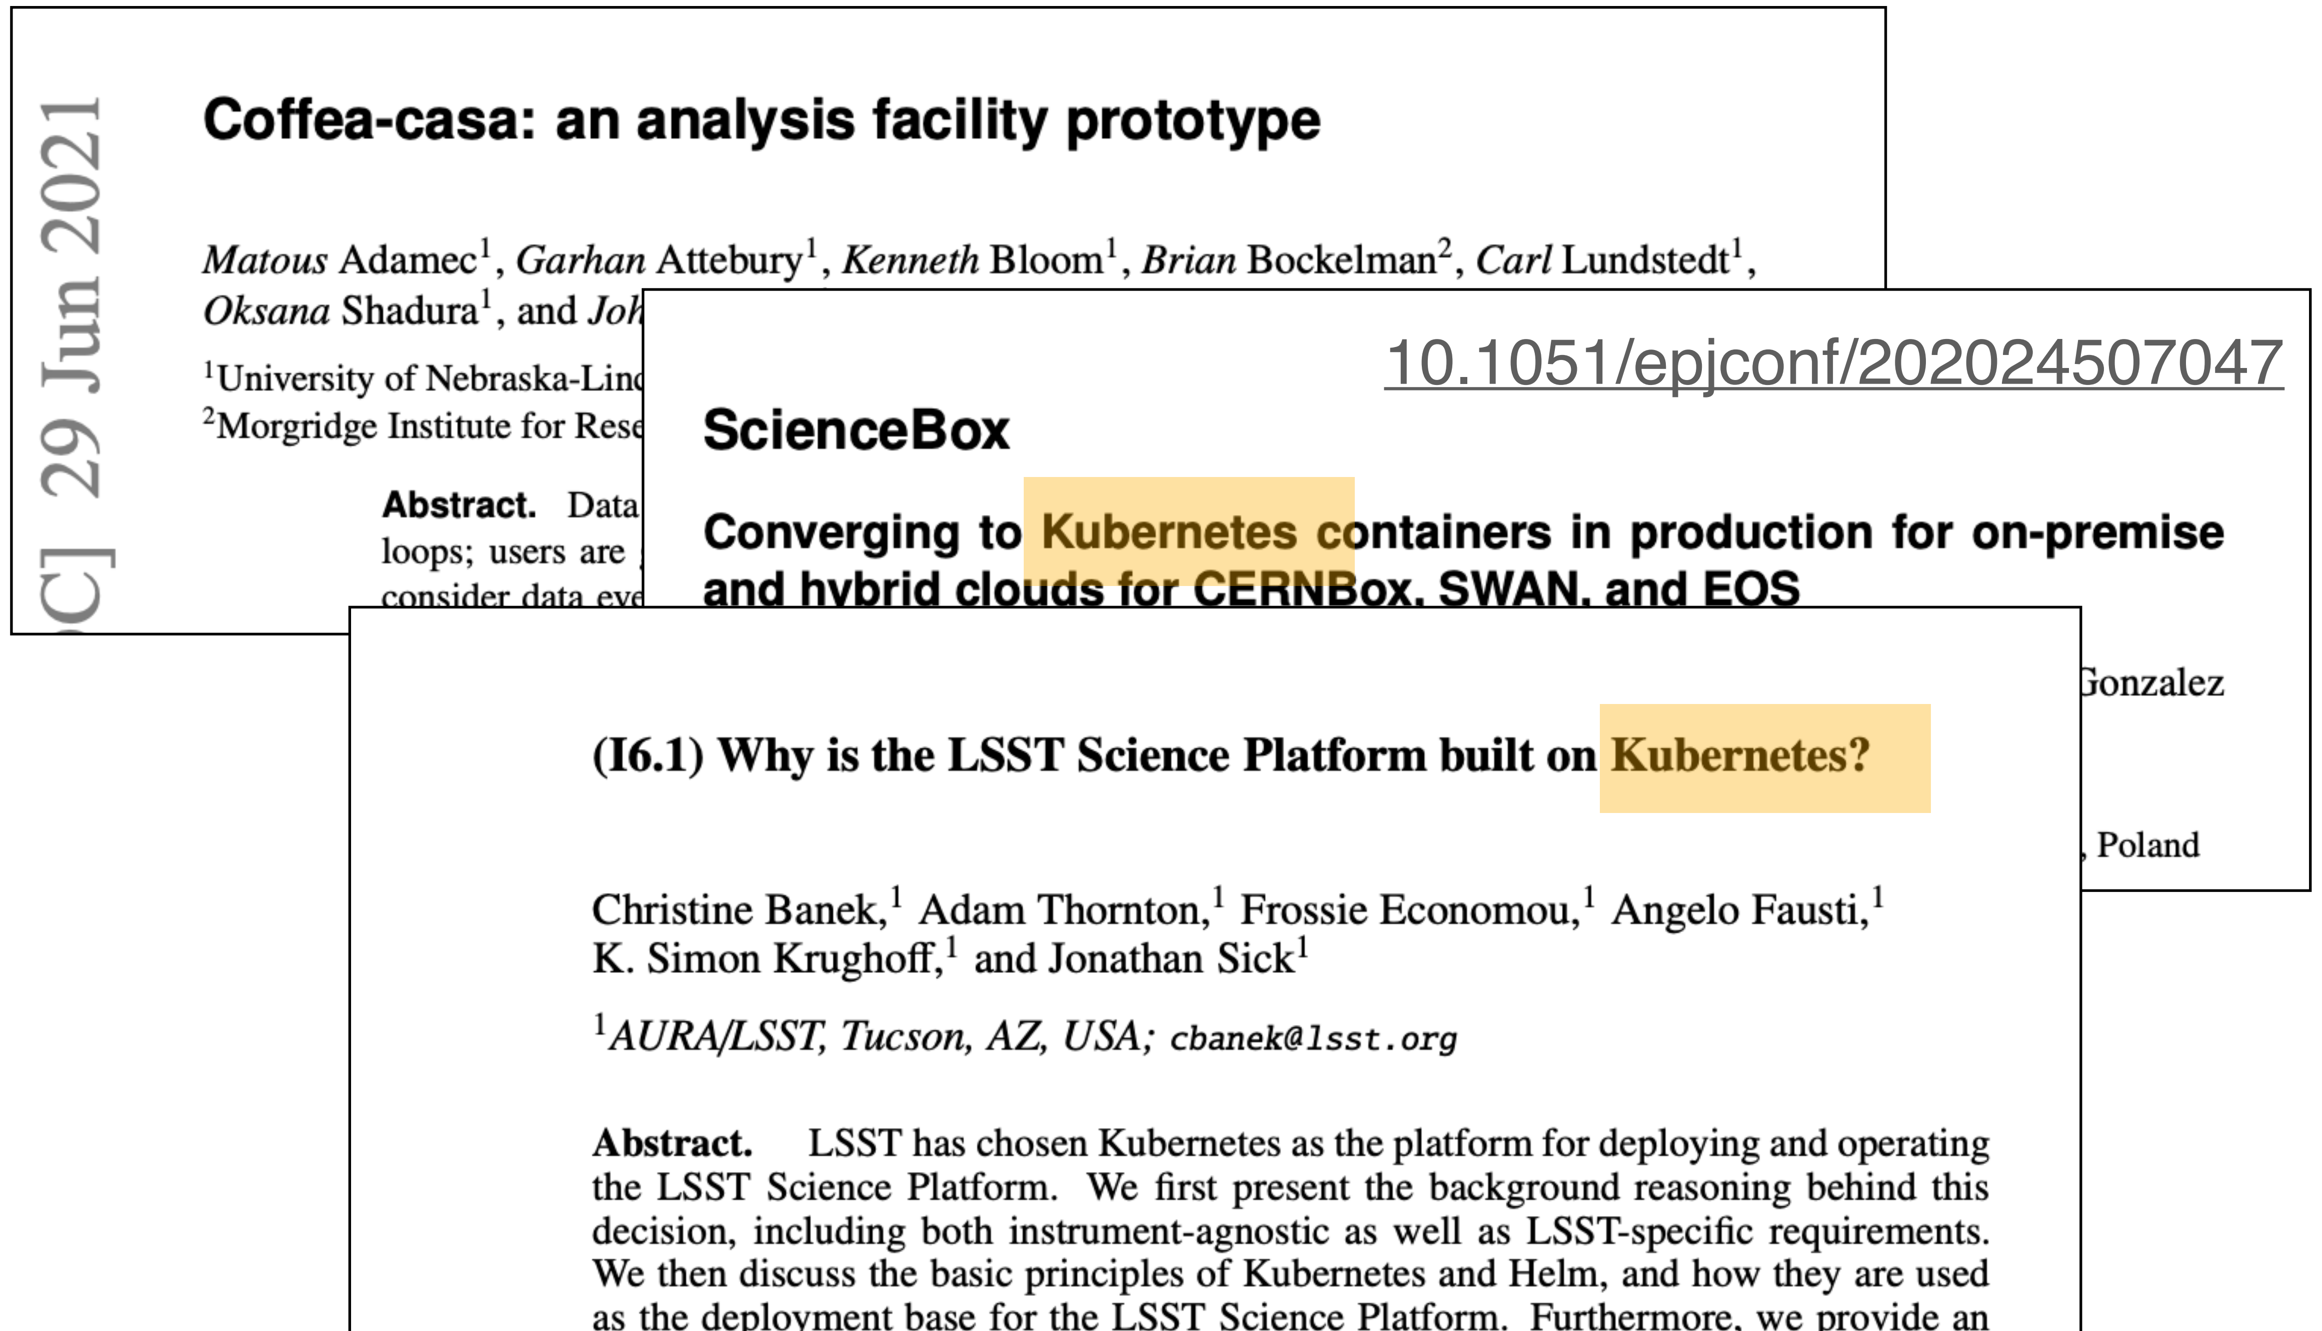

Bringing the Grid to the Cloud

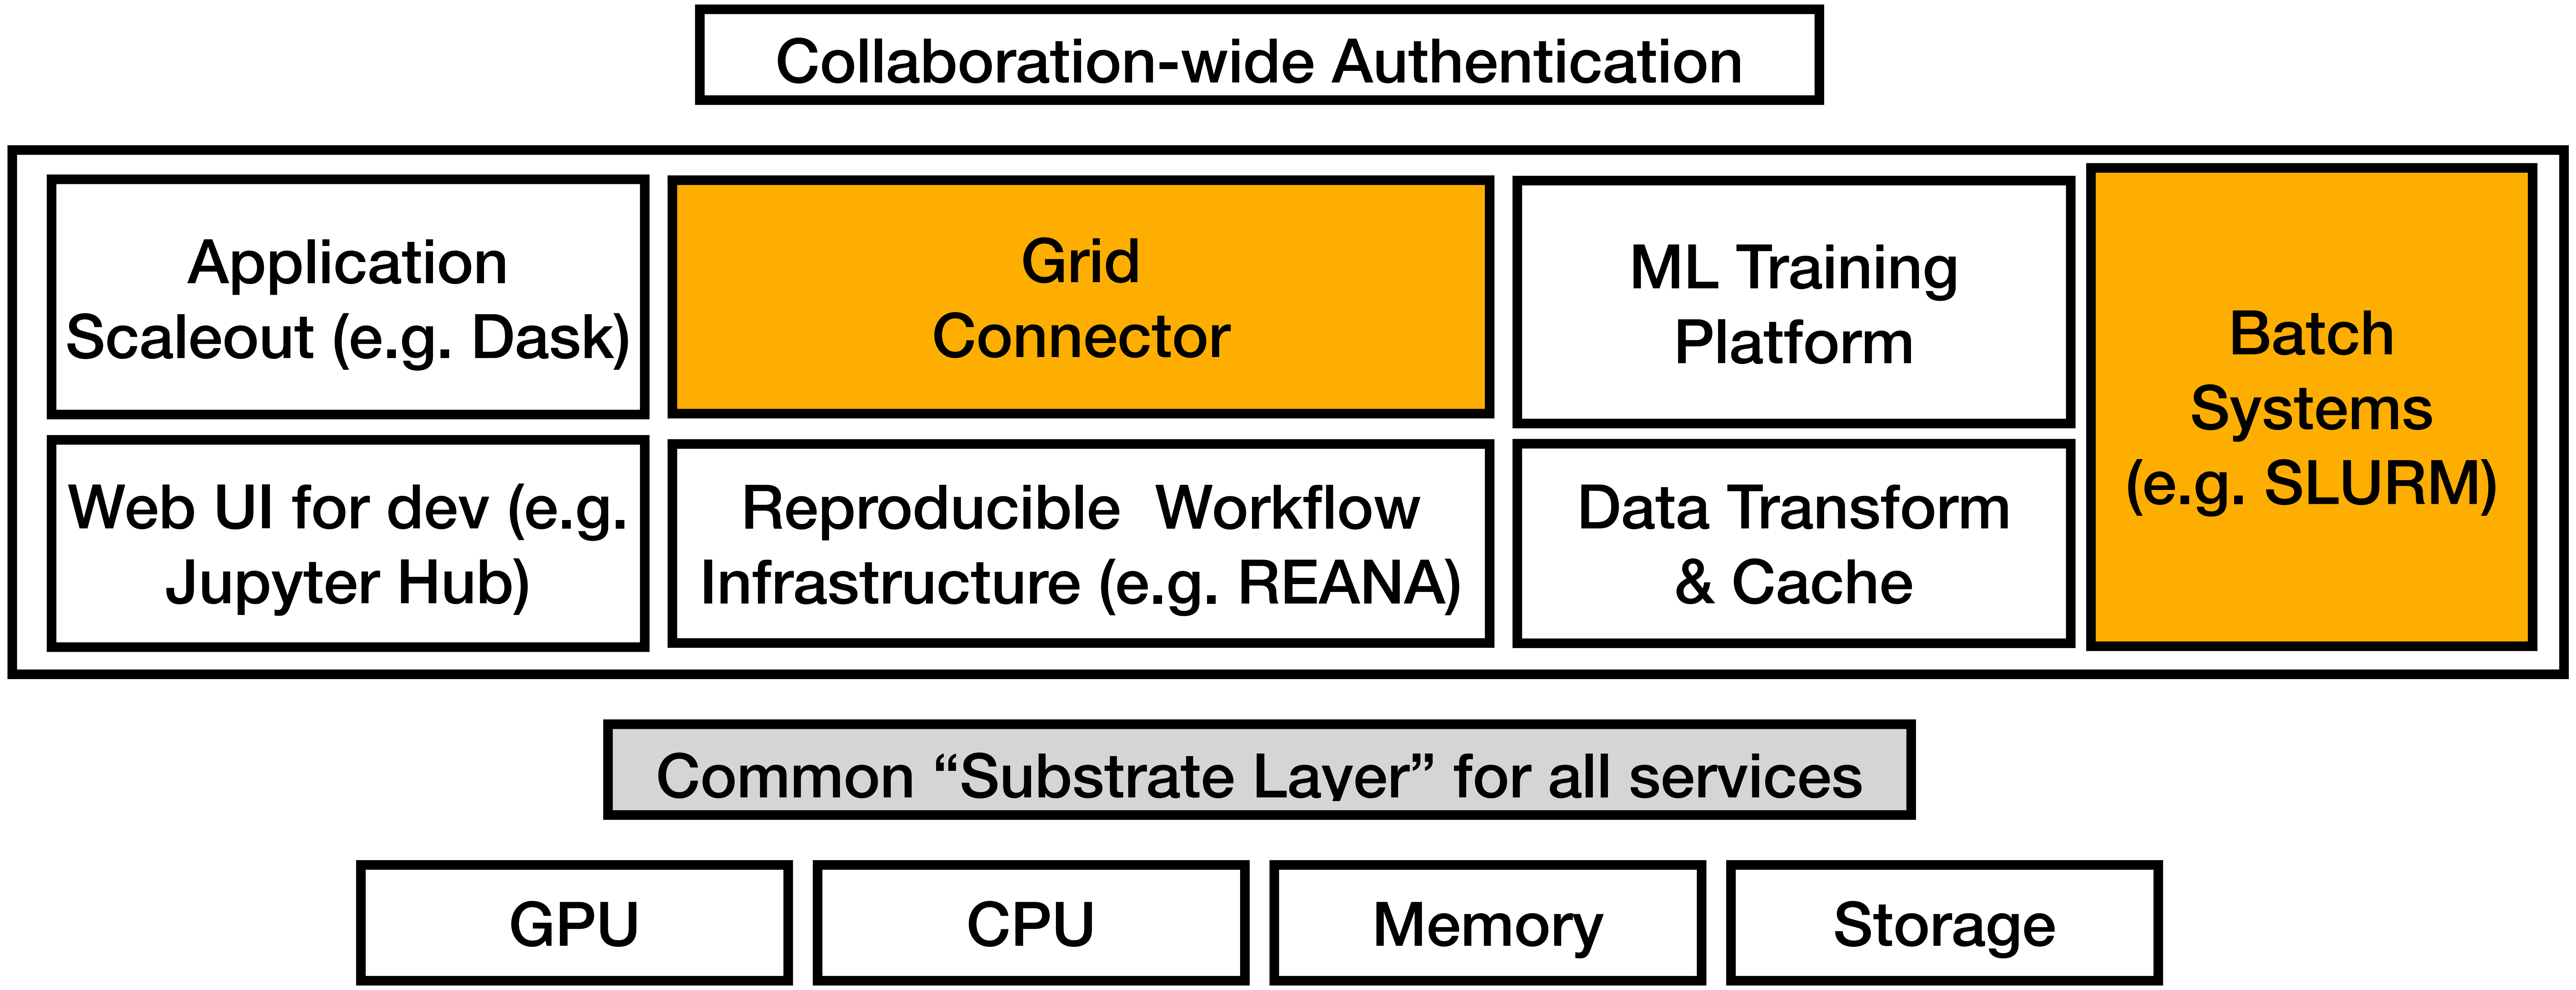

First things first: Grid Sites

Show we can deploy our existing tools on new infrastructure

Advantage: we were always set up to to integrate new systems.

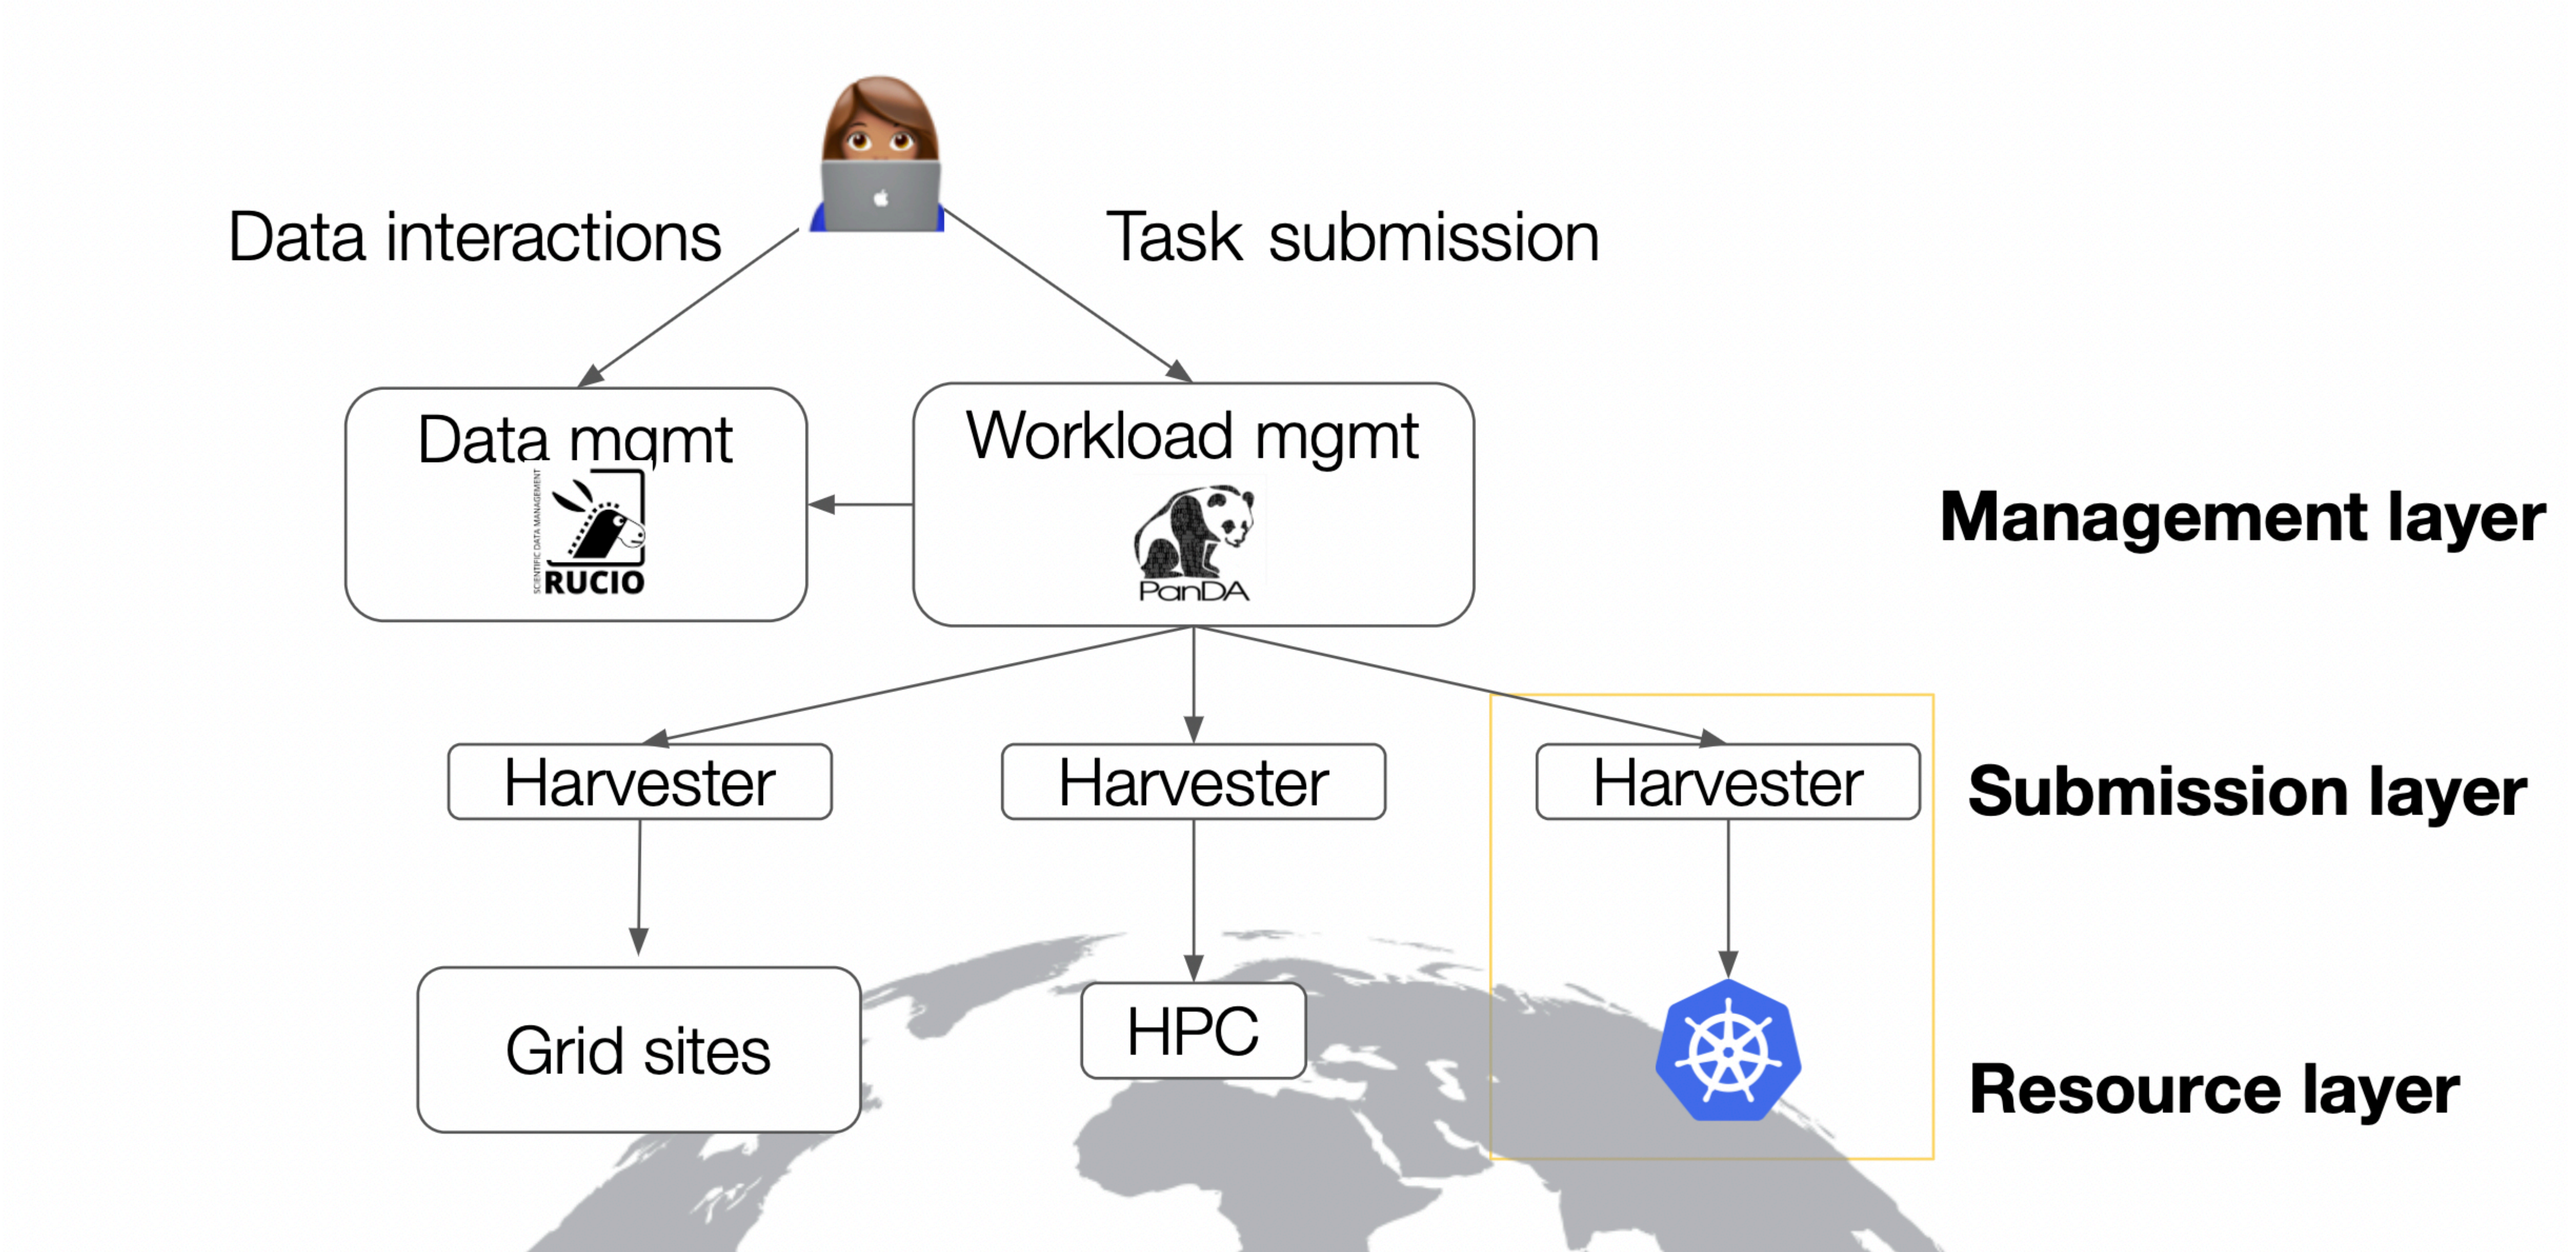

First things first: Grid Sites

Kubernetes as just another resource, receive grid jobs & run them

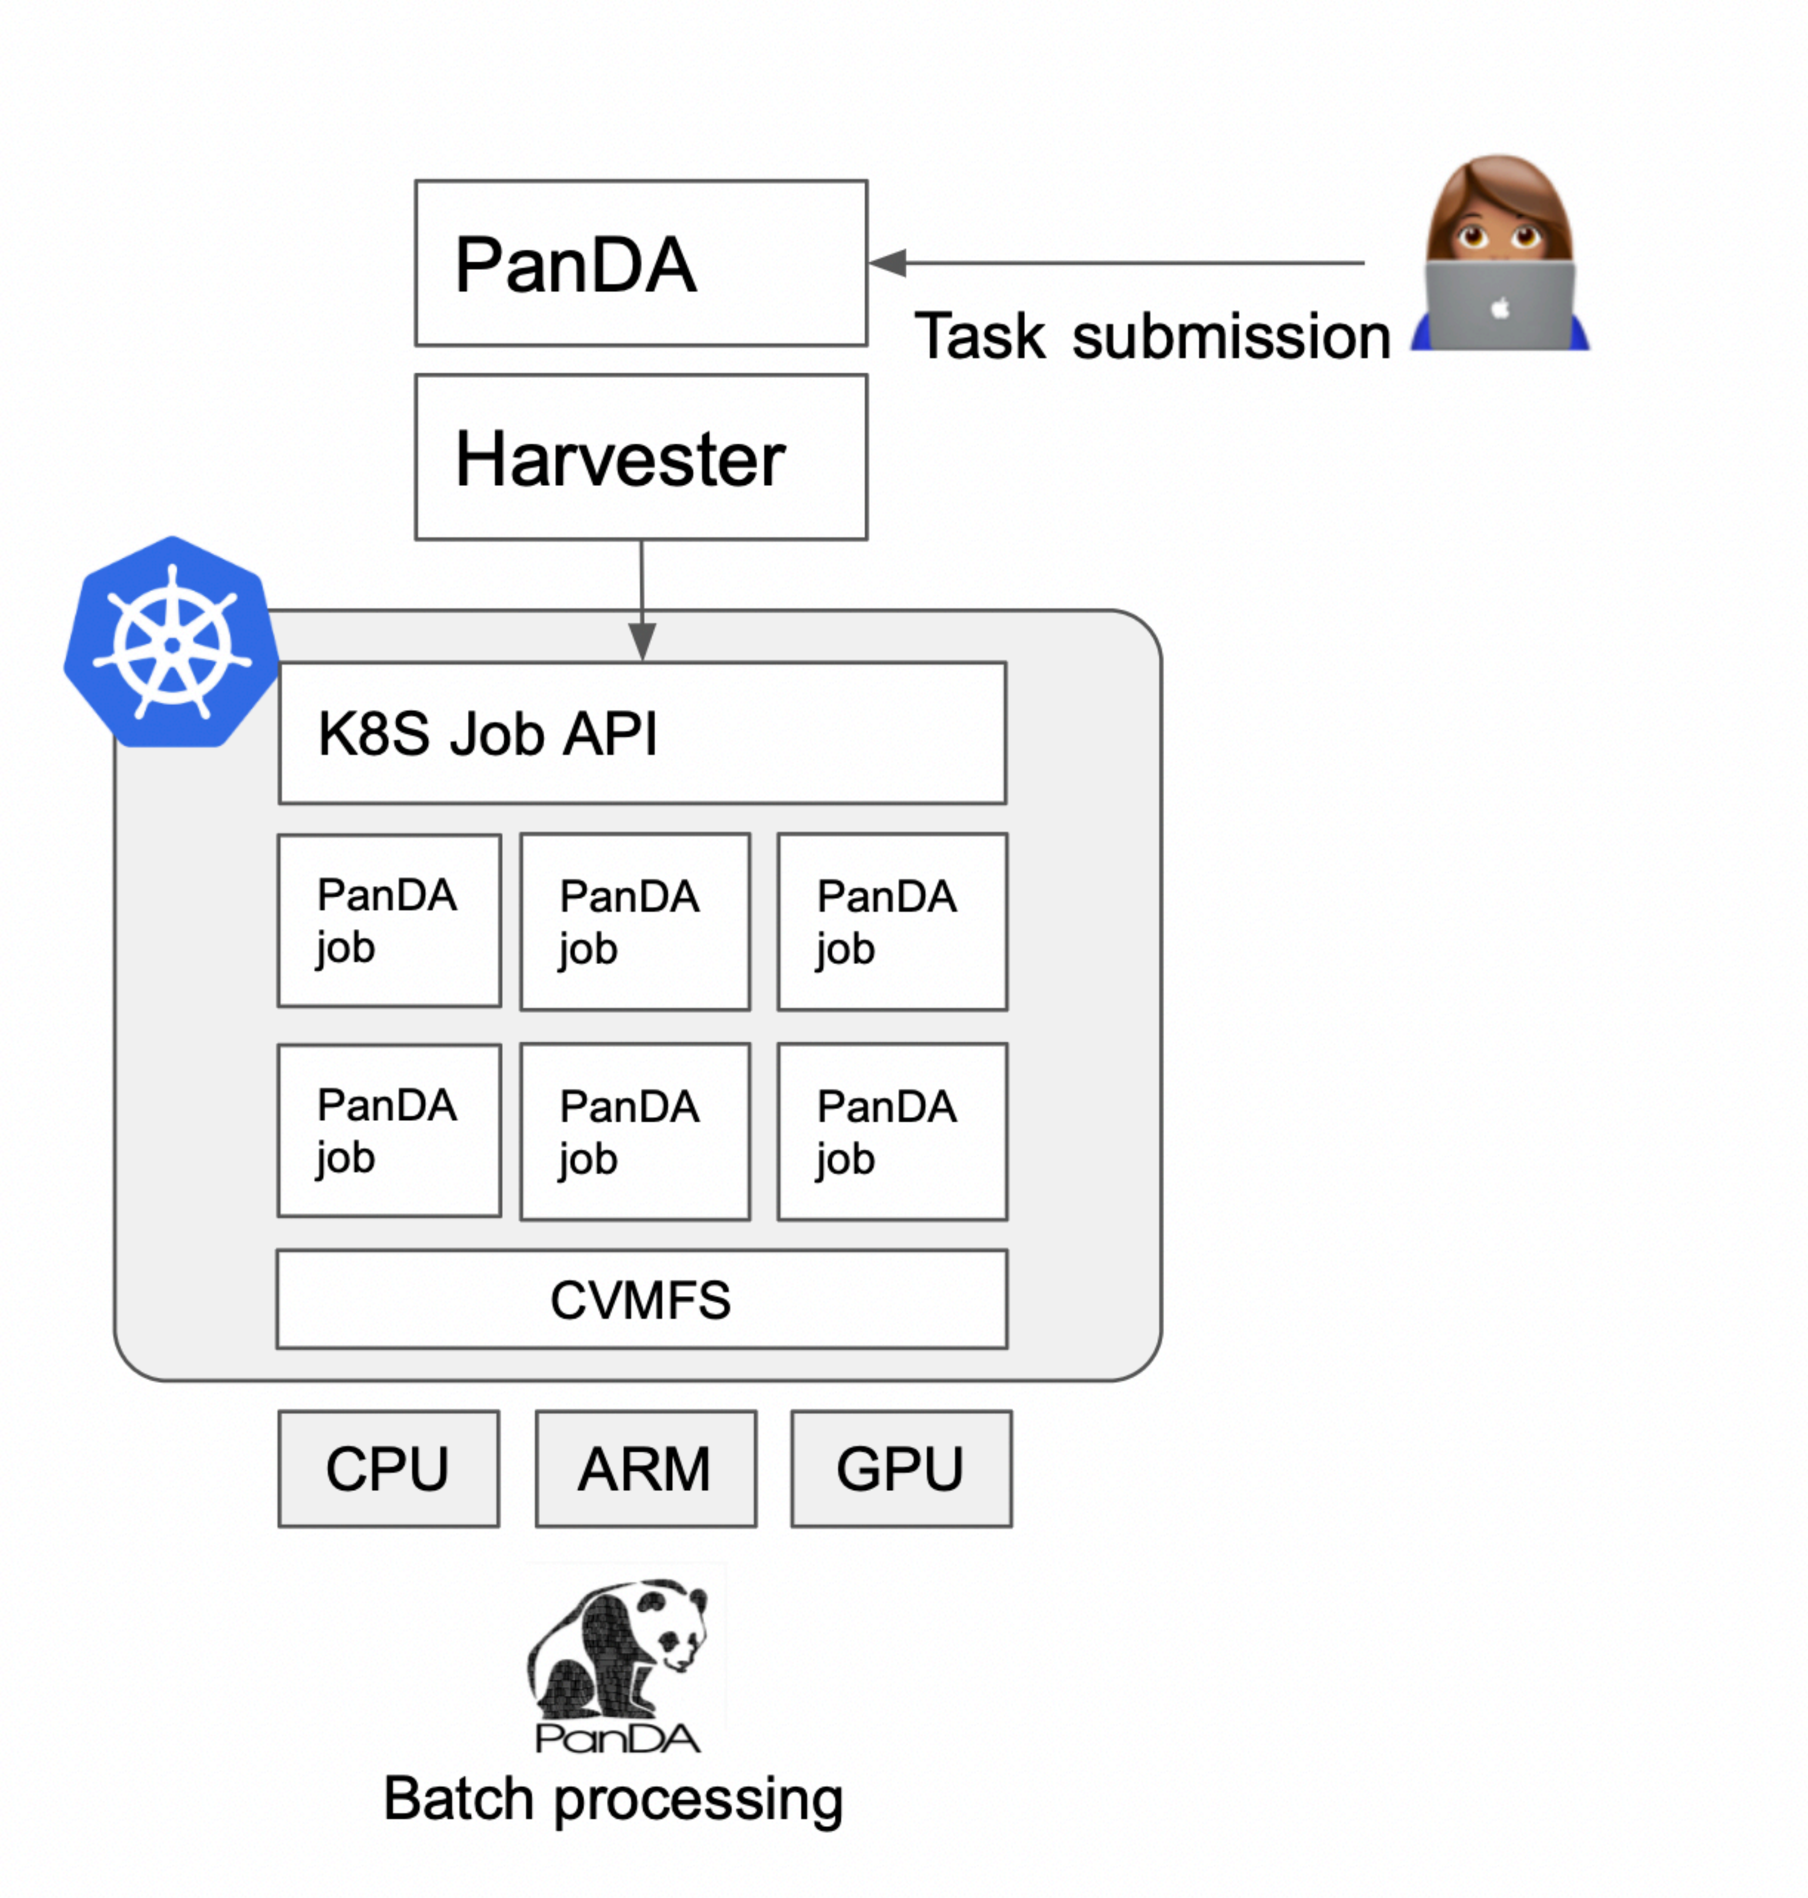

- no external batch system
- dynamic job placement across nodes to maximize utilization
- in cloud it's easy to add heterogeneous resources such as GPU or ARM

First things first: Grid Sites

Successfully running in production O(10k cores) across multiple institute cloud **and** clusters with the same Kubernetes-based setup.

Growing fraction of the grid. How far can we go with a single cluster?

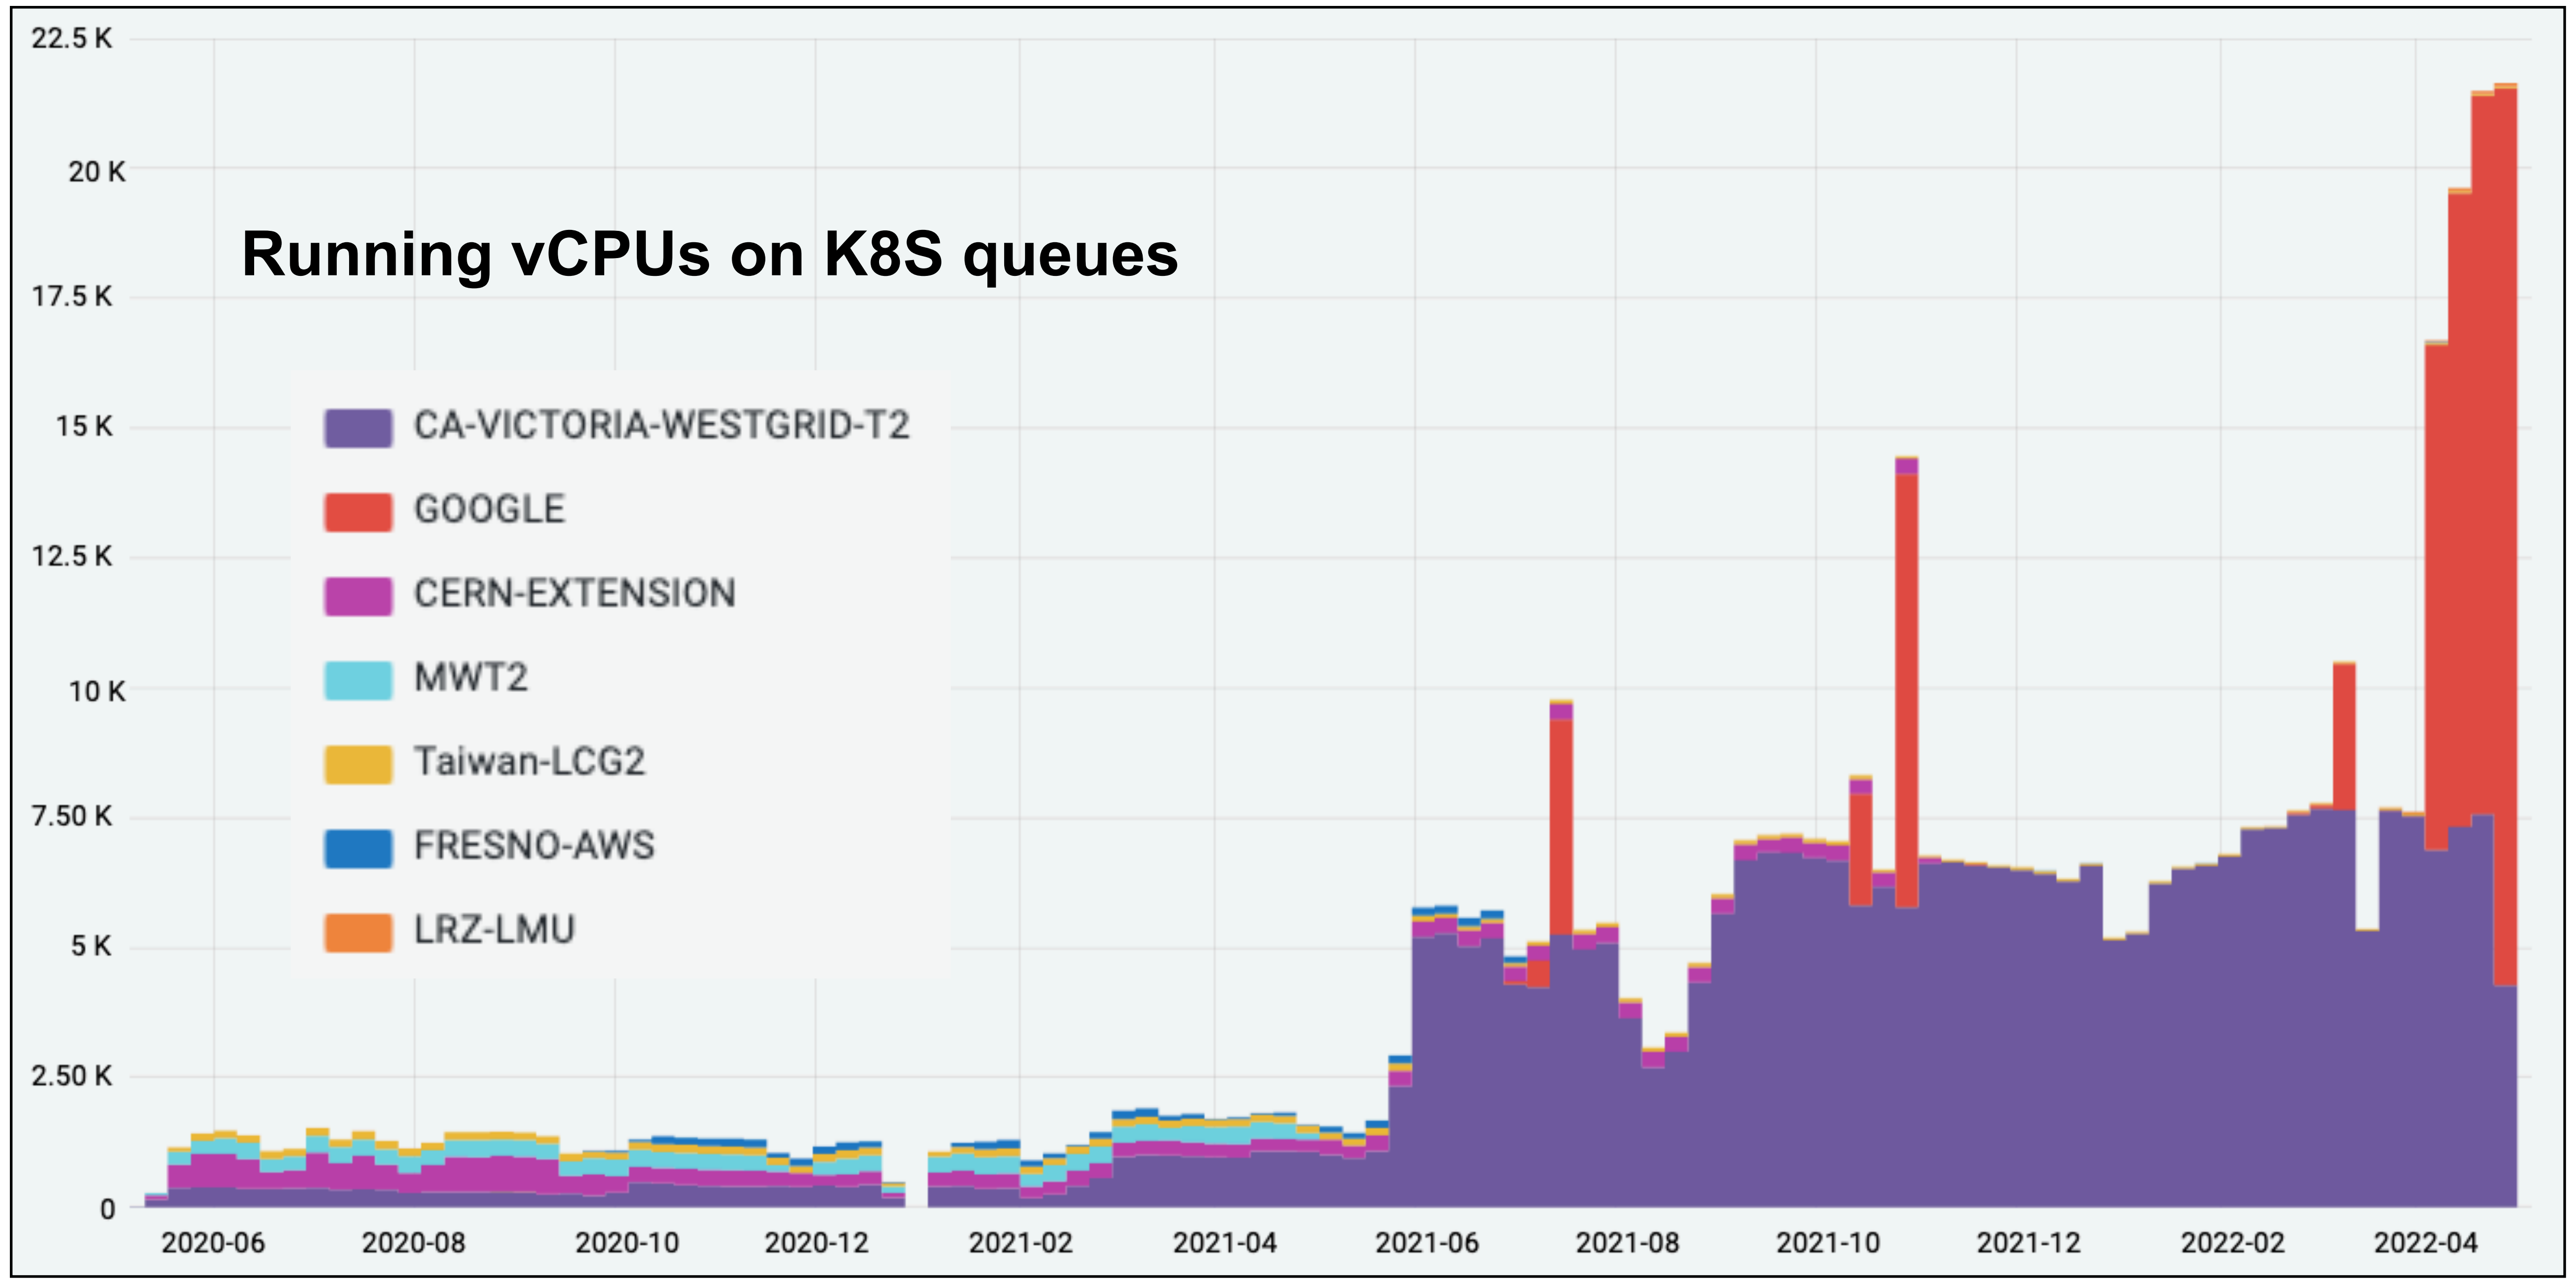

FRESNO STATE

Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften

THE UNIVERSITY OF CHICAGO

aws

CERN

University of Victoria

Google Cloud

Cloud-specifics

In order to seriously scale & benefit from cloud elasticity a few adjustments to the setup

- Only a single small node is on 24/7
- HEP is embarrassingly parallel can just resubmit job → use pre-emptible VMs 5x cheaper
- to scale use “big nodes” e.g. 80 core: less room for failure (e.g. mounting fileystems)

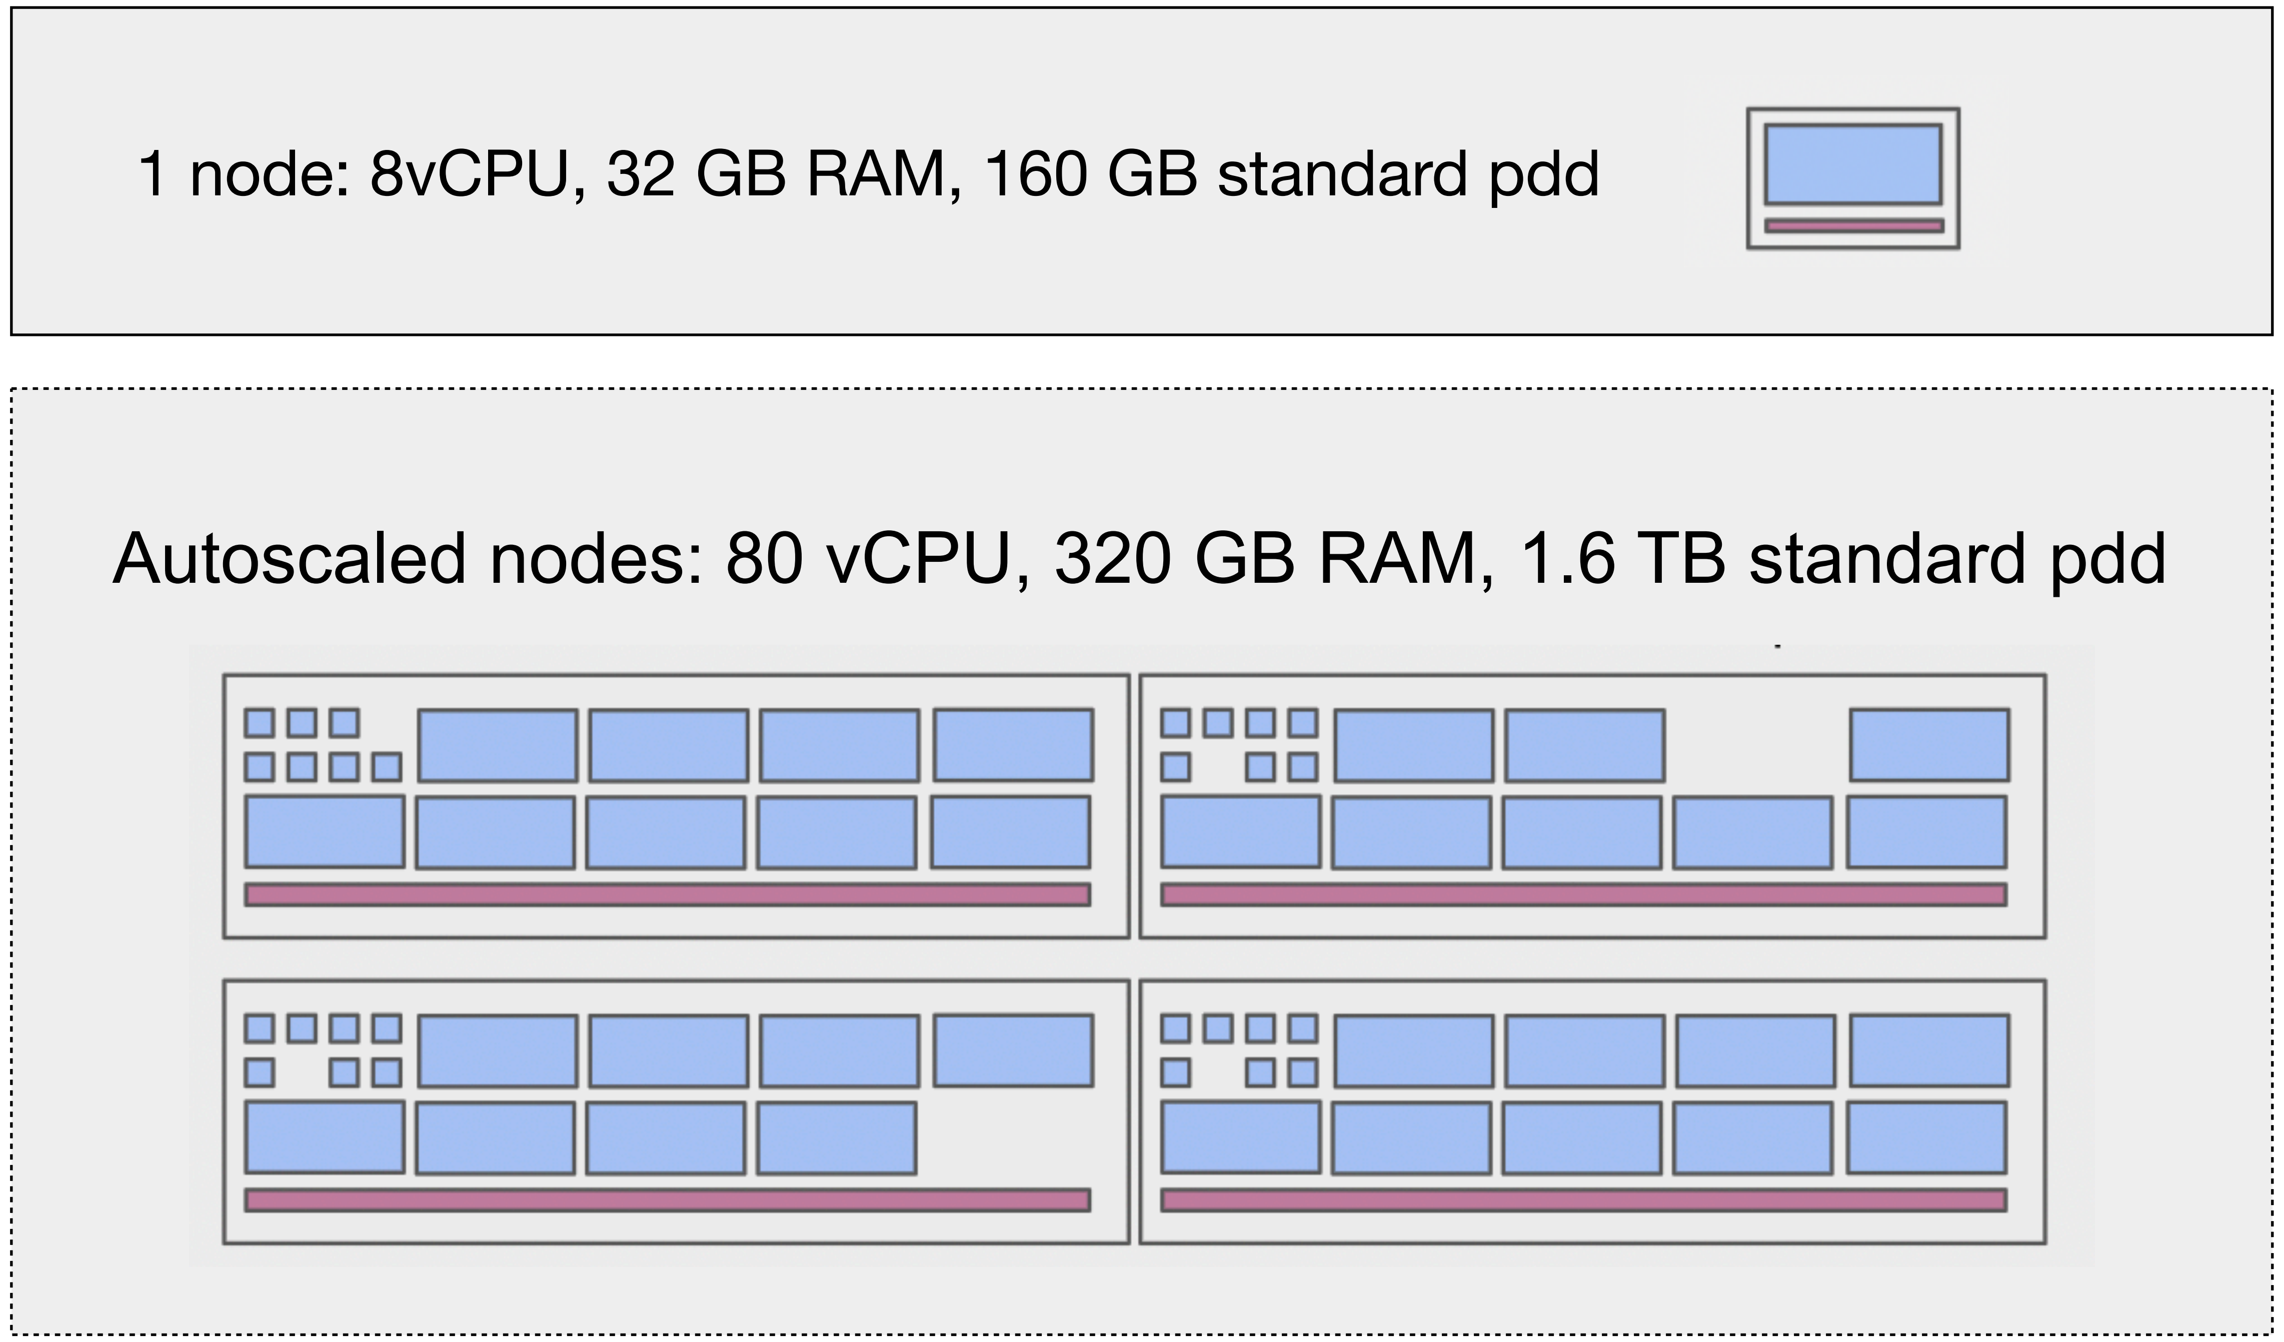

Scaling Up on the Cloud

With adjustments in place, we can comfortably scale to 100k cores on a single Kubernetes cluster running standard Grid workflows.

- 10k vs 100k makes little difference
- Only 1% job failure rate even on pre-emptible VMs

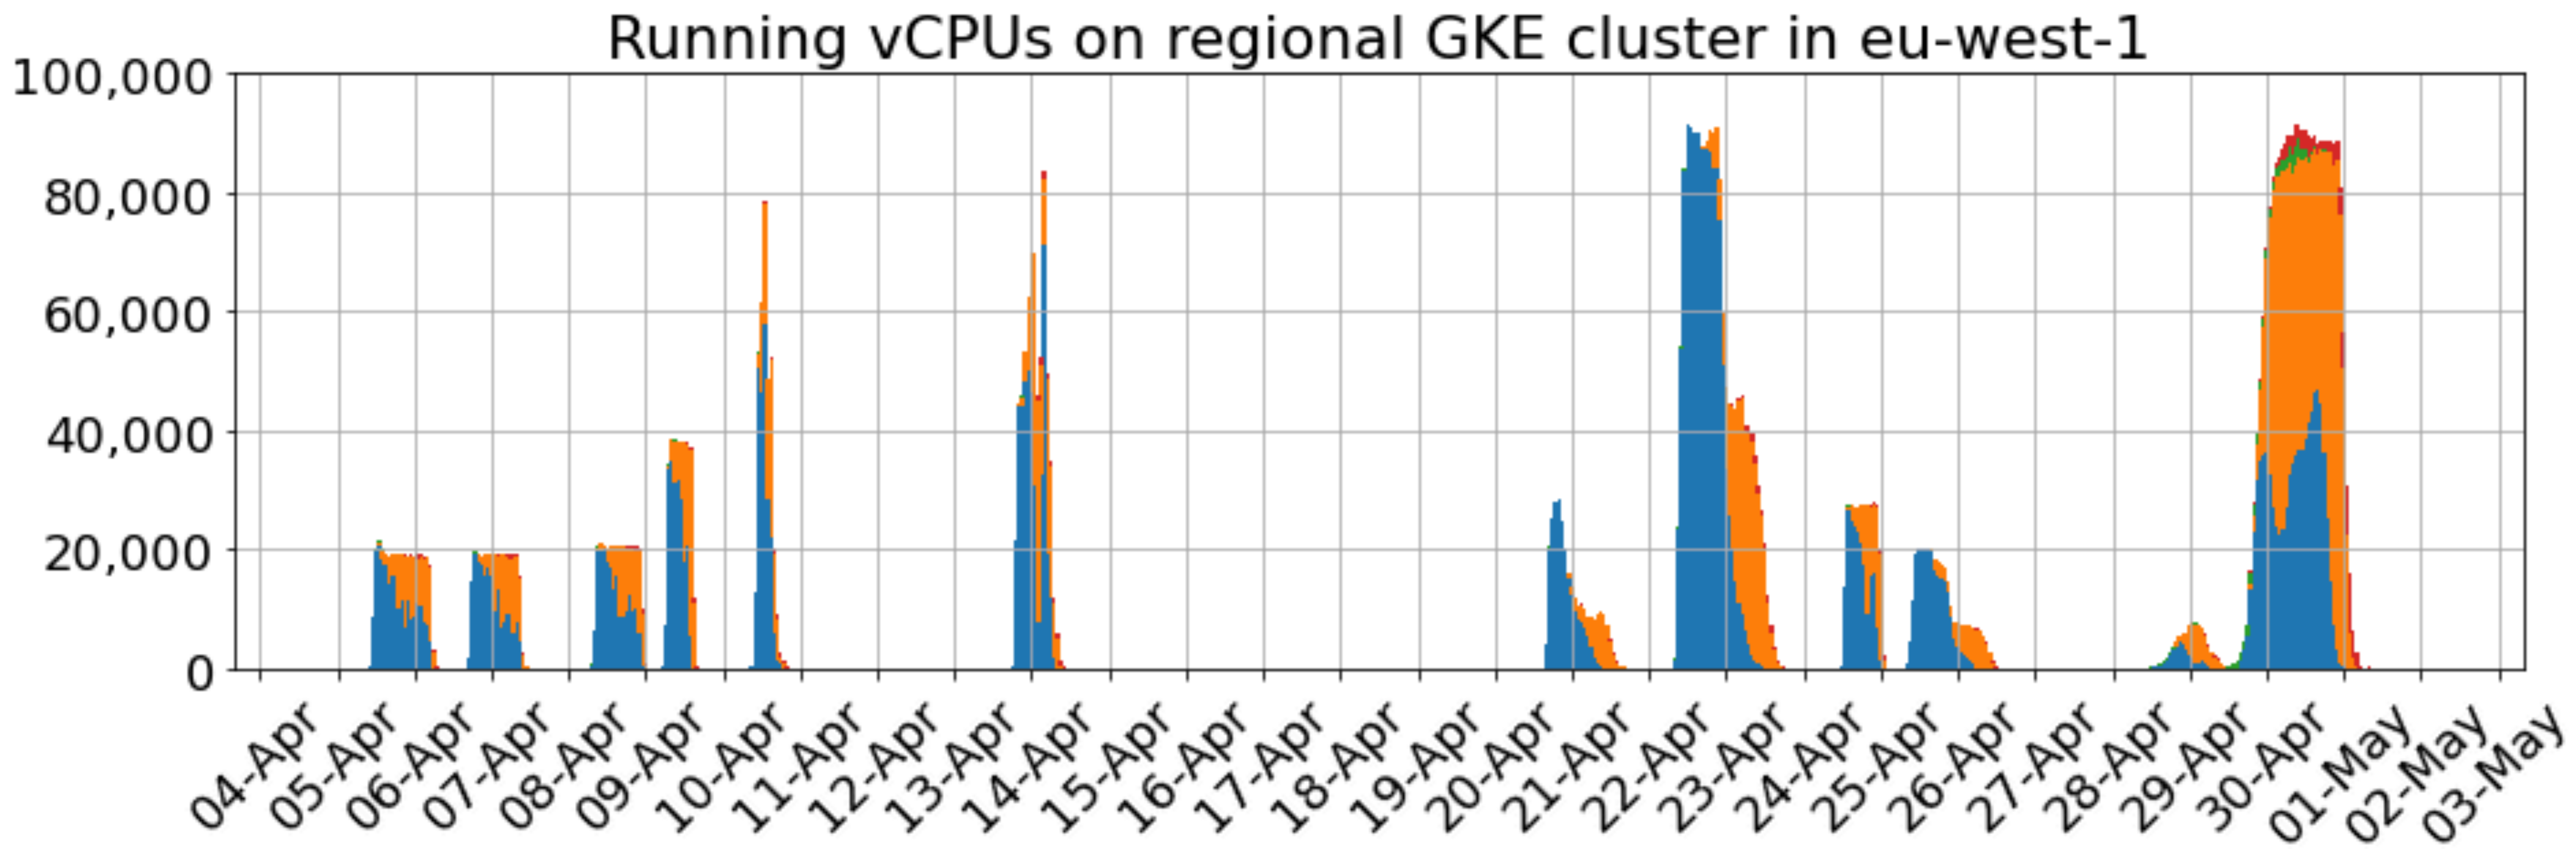

With cloud resources we could progressively scale: 10k → 20k → 40k → 100k cores

Scaling Up on the Cloud

At 100k cores, a single cluster is a sizable fraction to the overall grid

- should be able to scale more, but started to see instabilities > 100k

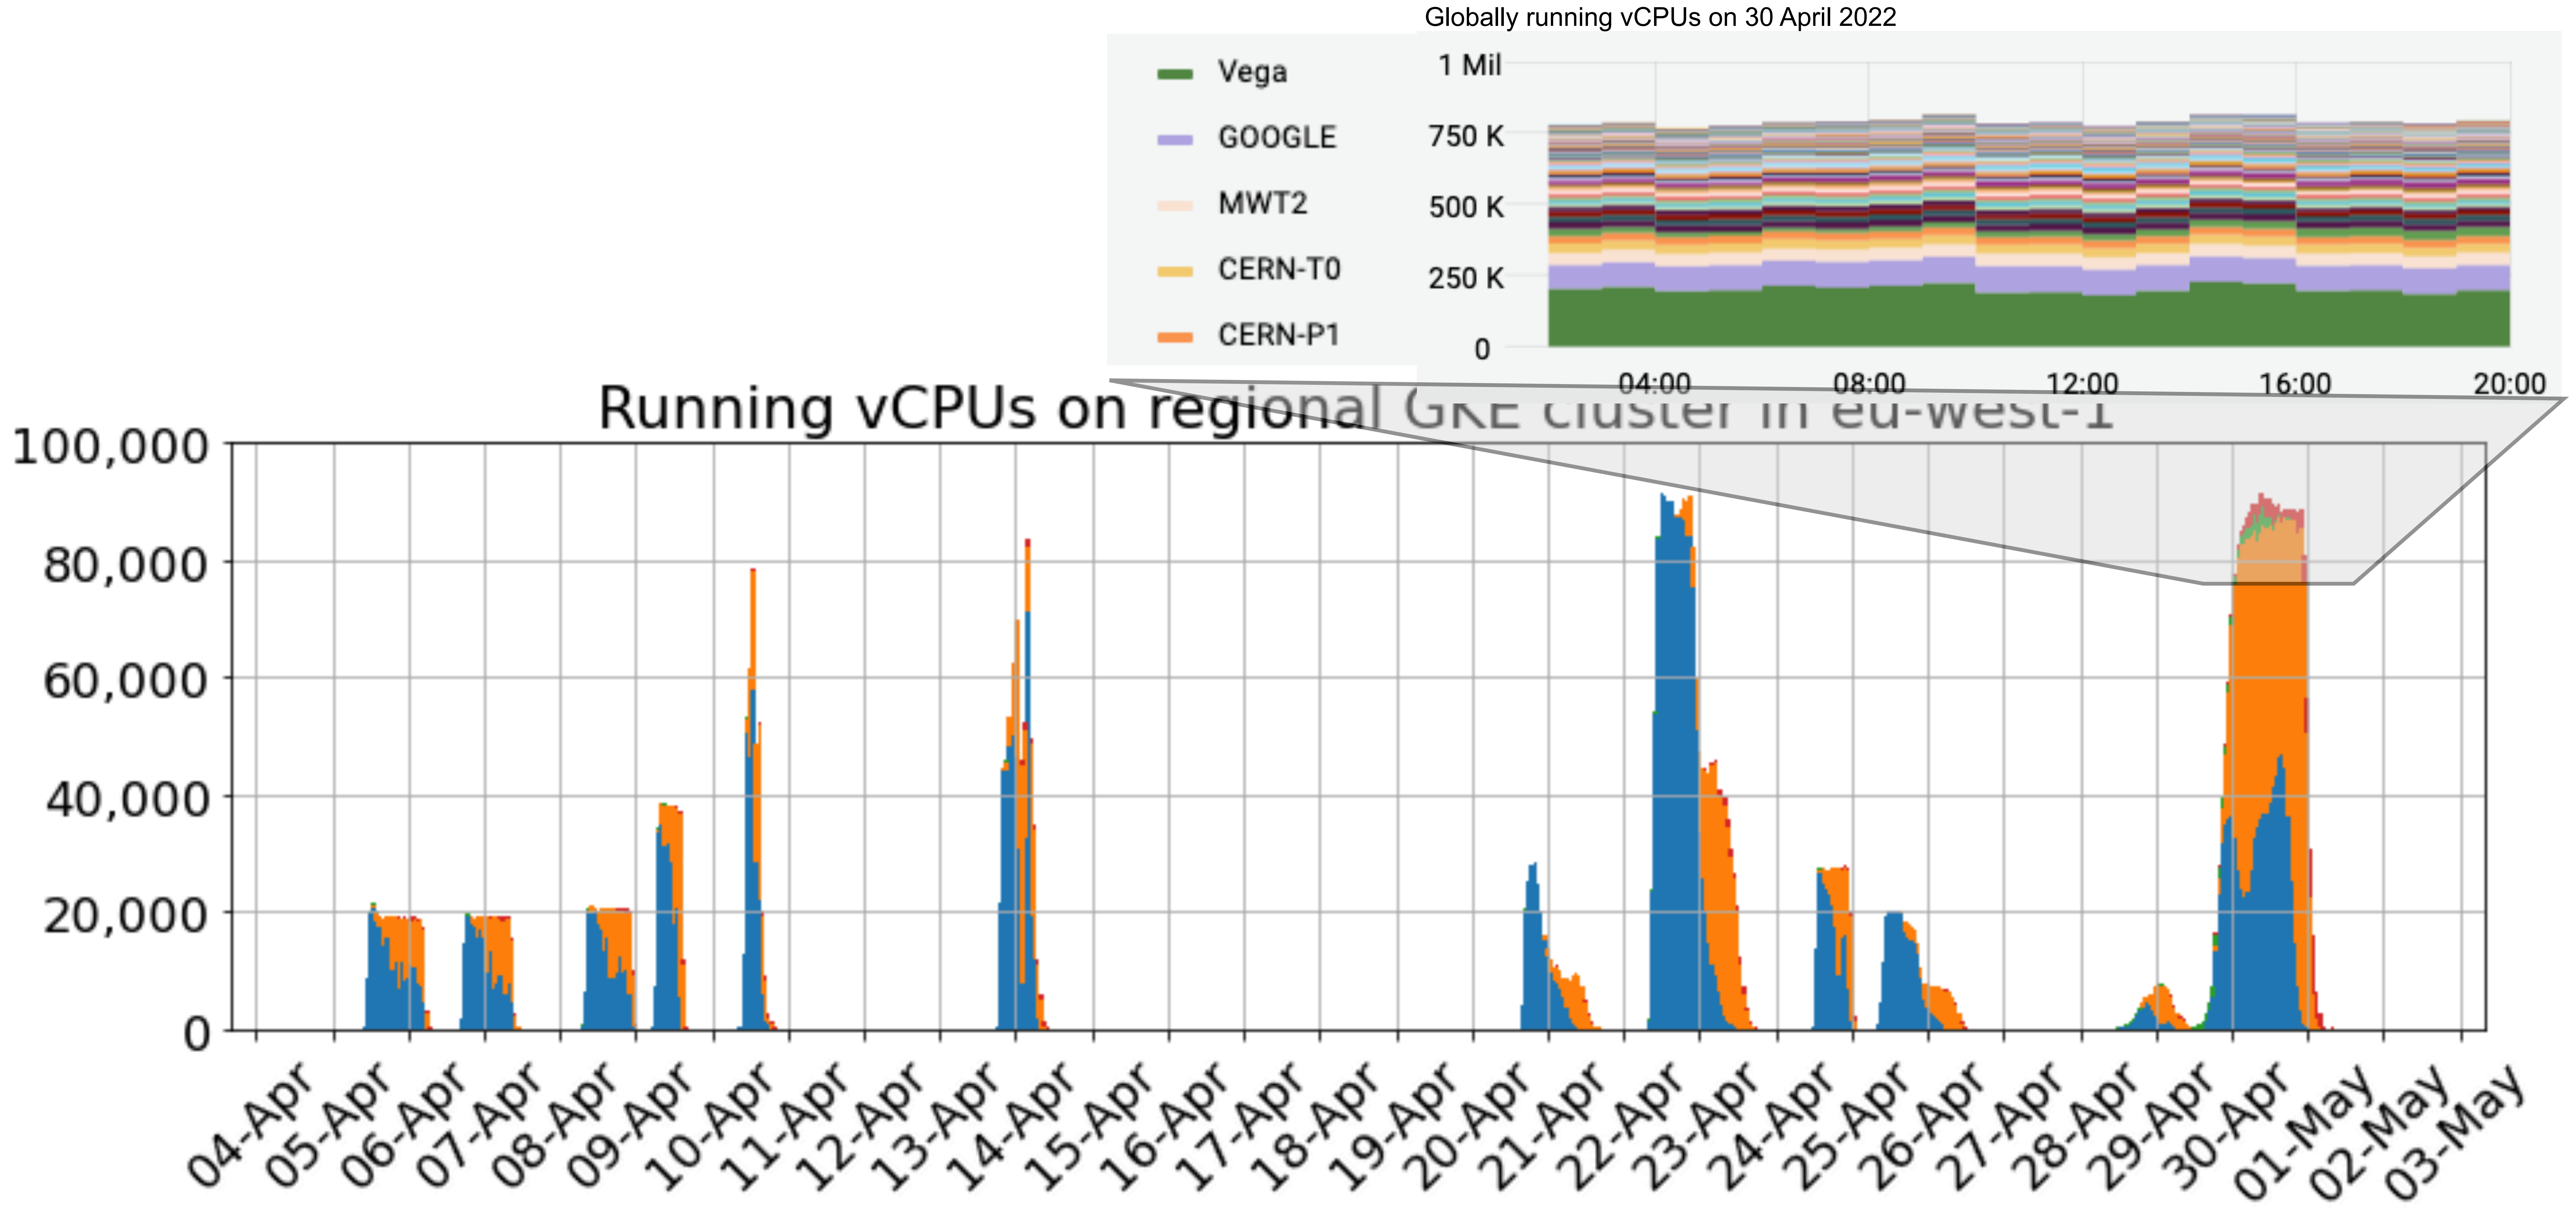

$O(10^8)$ events

$O(10^5)$ vCPUs

$O(10^4)$ Pods

$O(10^3)$ Nodes

$O(1)$ day

$O(1)$ Harvester instance

<1 Engineer

Beyond Simple Batch

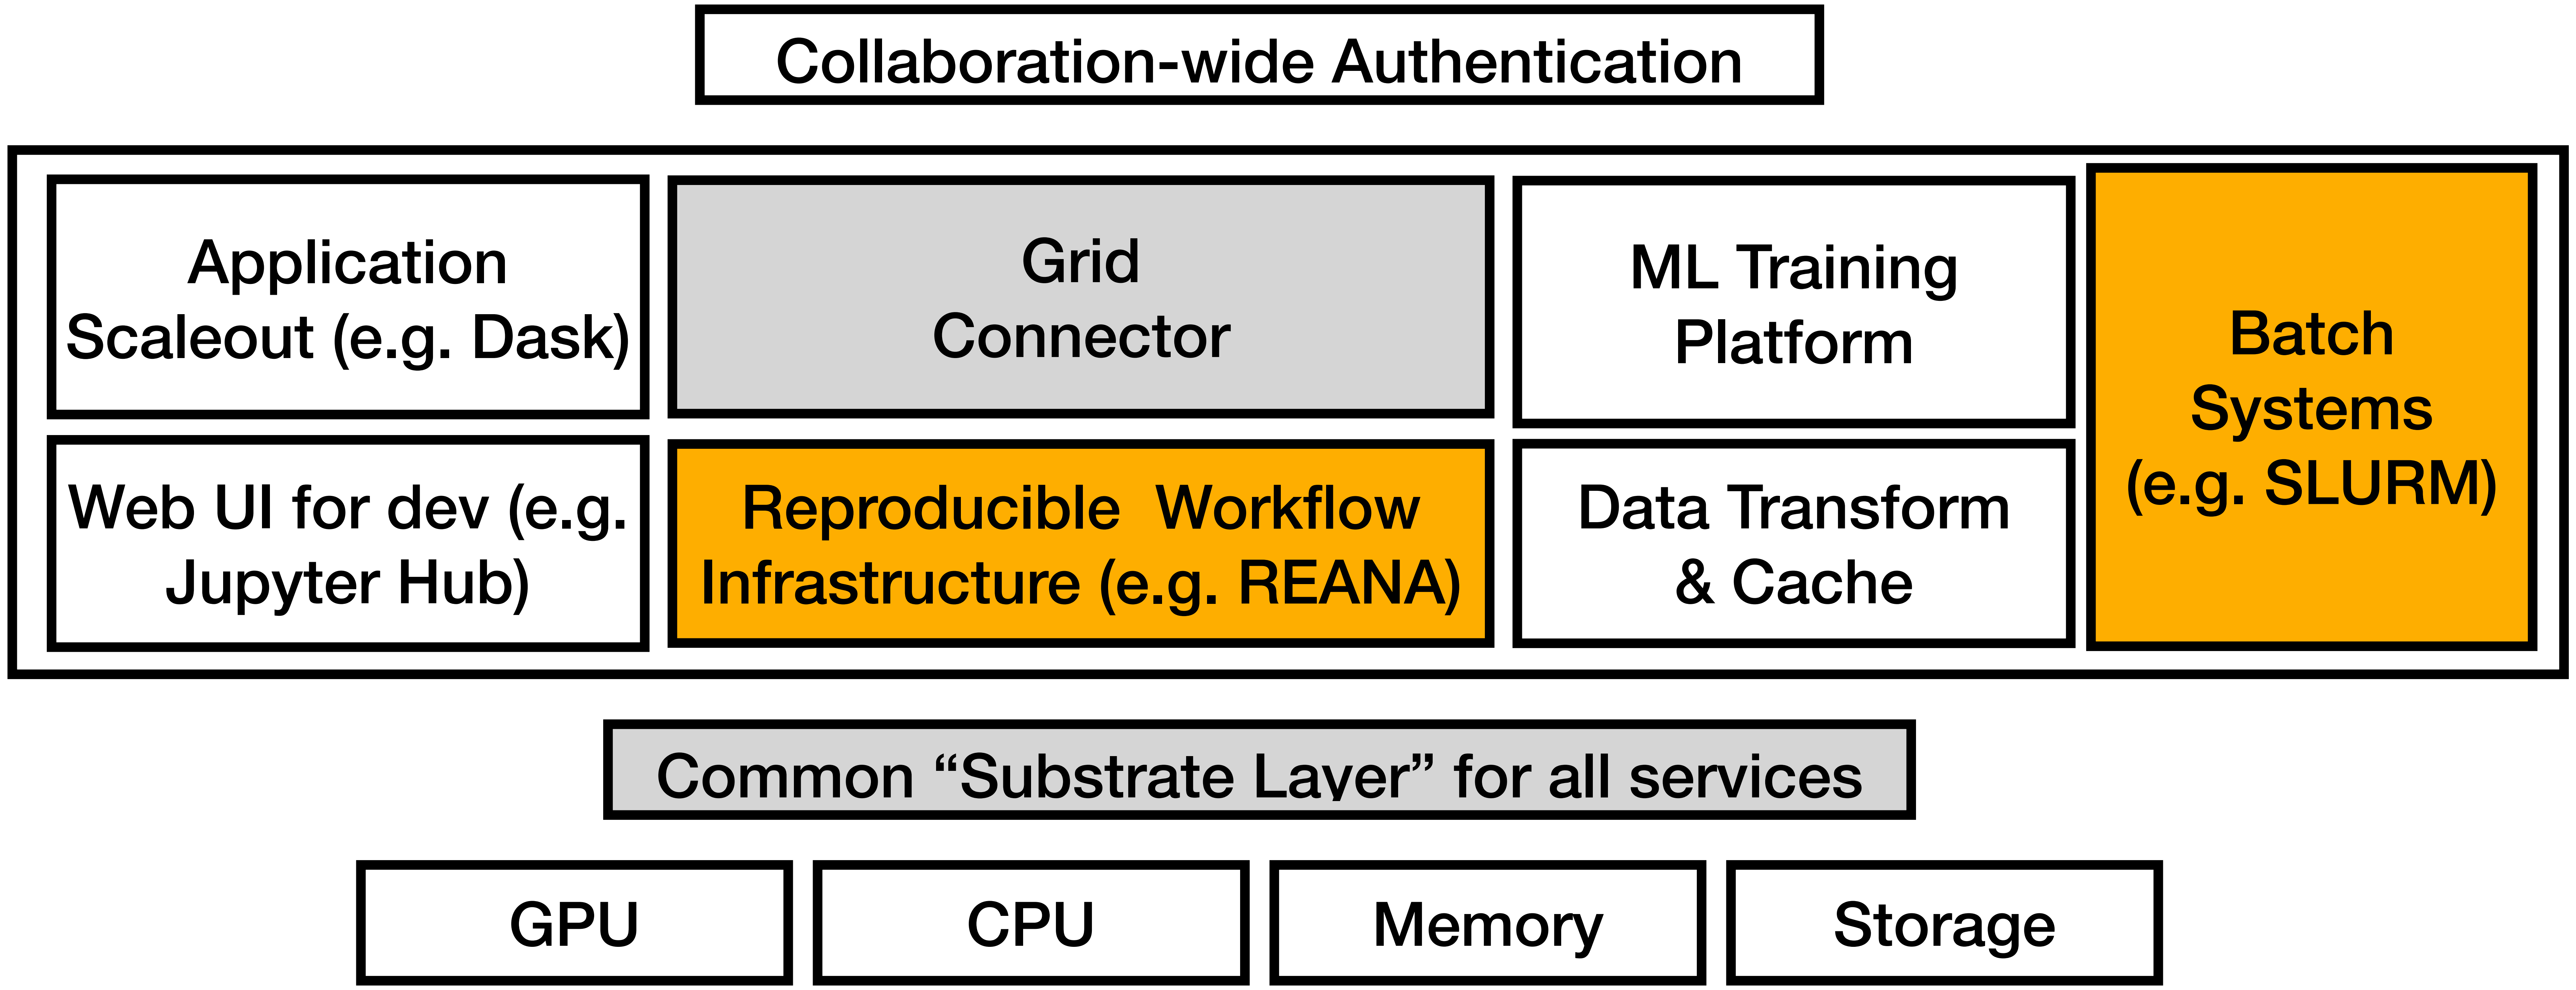

Adding Reproducible Workflows

Grid as it is today very optimized for "collaboration-wide" tasks (simulation and reconstruction) very uniform data processing

User analysis is much more heterogeneous: complex workflows, each analysis unique: ***can we map this onto the same infrastructure?***

- opportunity to expand what we mean by "the grid"

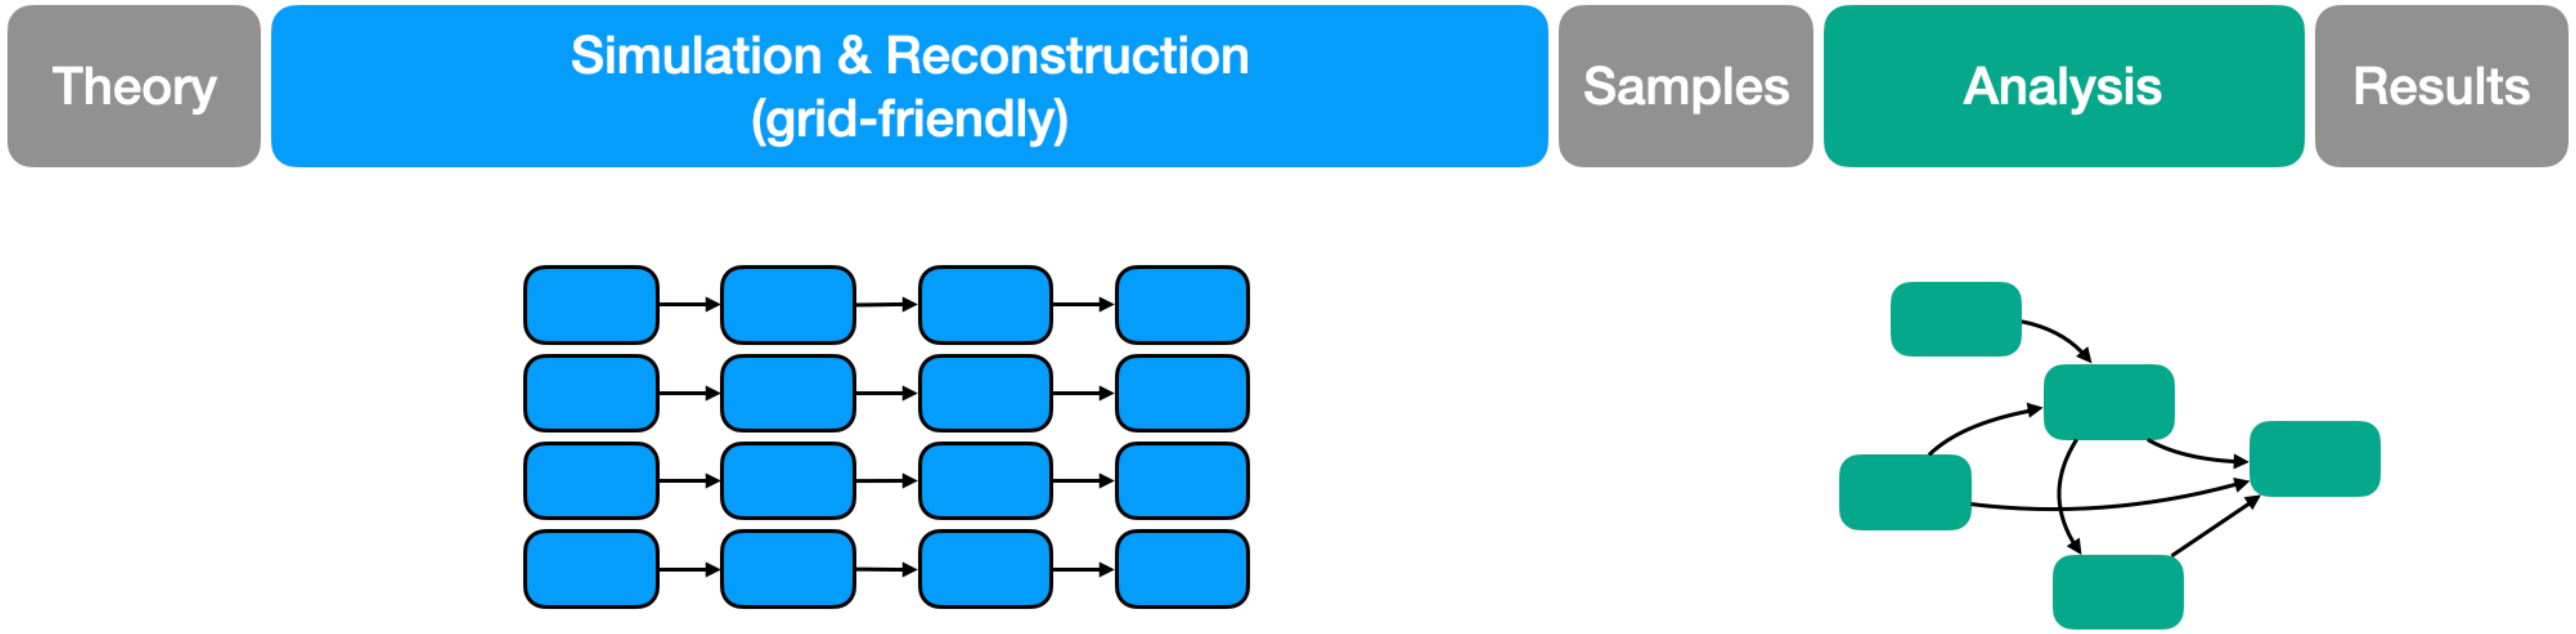

REANA

Platform for resuable (the “R” in FAIR) analysis workflows on cloud infrastructure.

Idea: describe analysis as a workflow of containerized tasks, let workflow system handle specifics of submission

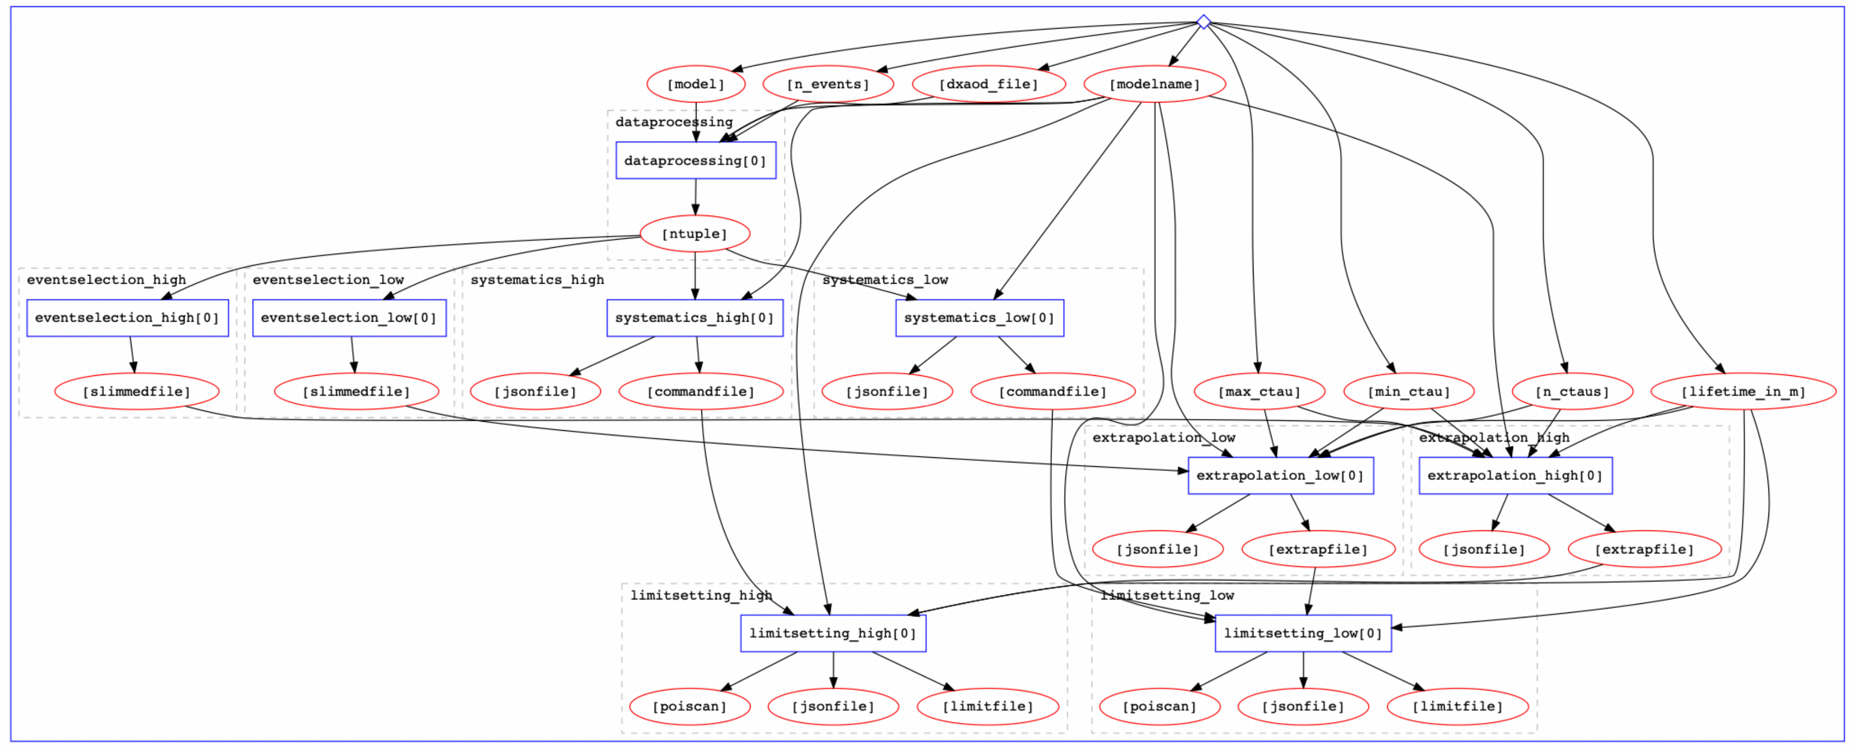

Many workflow systems in use, REANA aims to support many of them on a unified level

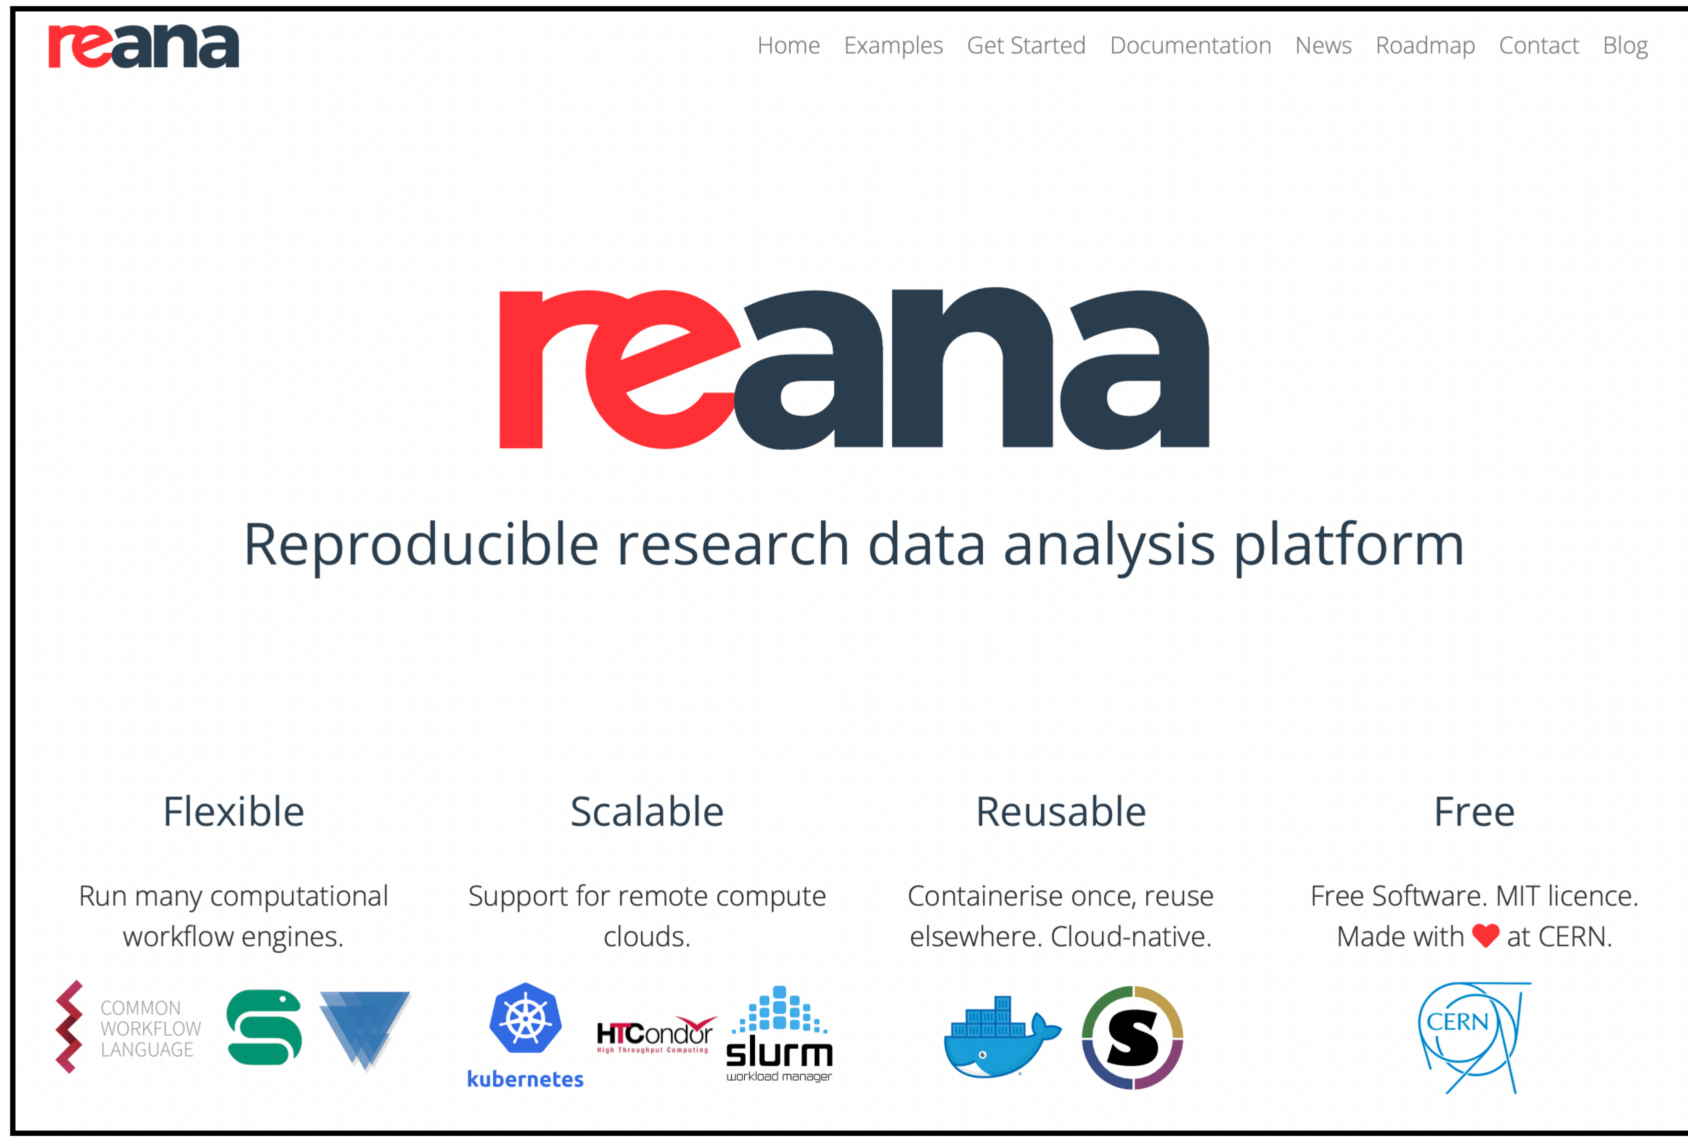

Key Use Case: RECAST

Idea: most analyses are sensitive to more physics than they originally targeted. To get the most physics out, we should “reuse” them:

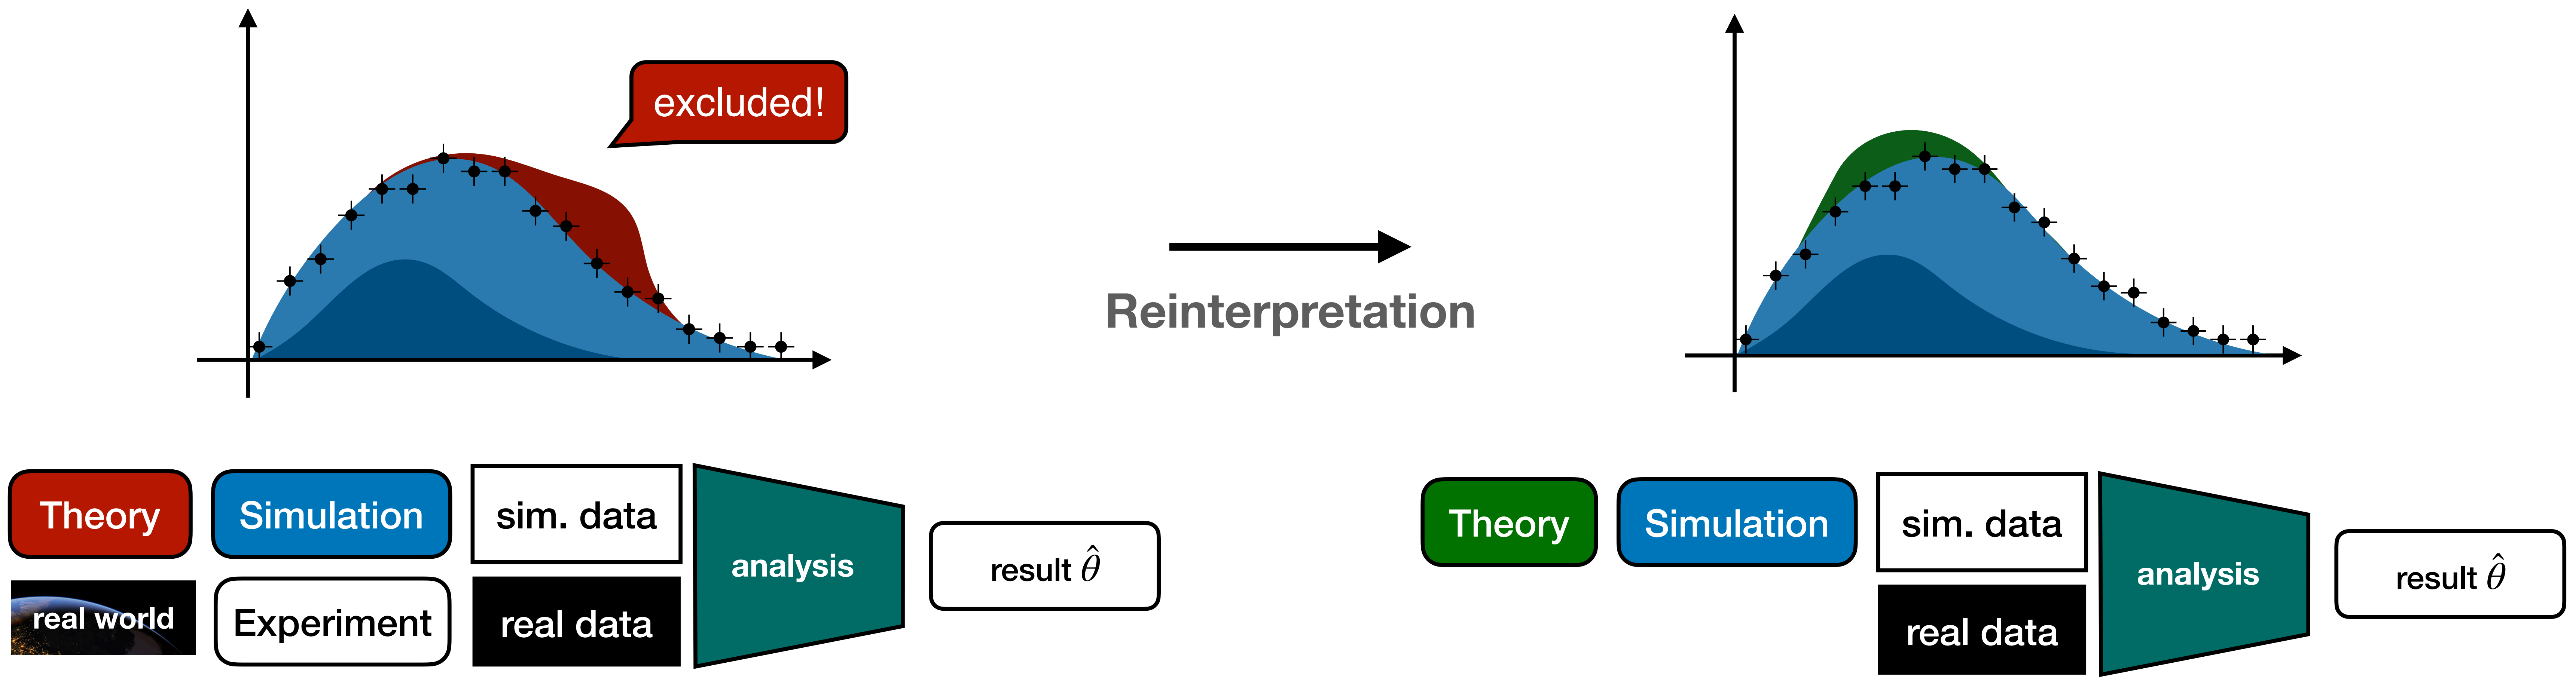

For this we need ability to

- preserve analyses as re-suable objects → containerized workflows
- catalogue them (i.e. findable - the “F” in FAIR)
- re-run them on common infrastructure

Examples of Reinterpretations

Starting to get new physics results from these efforts for “hand-picked” analyses and individual new theories.

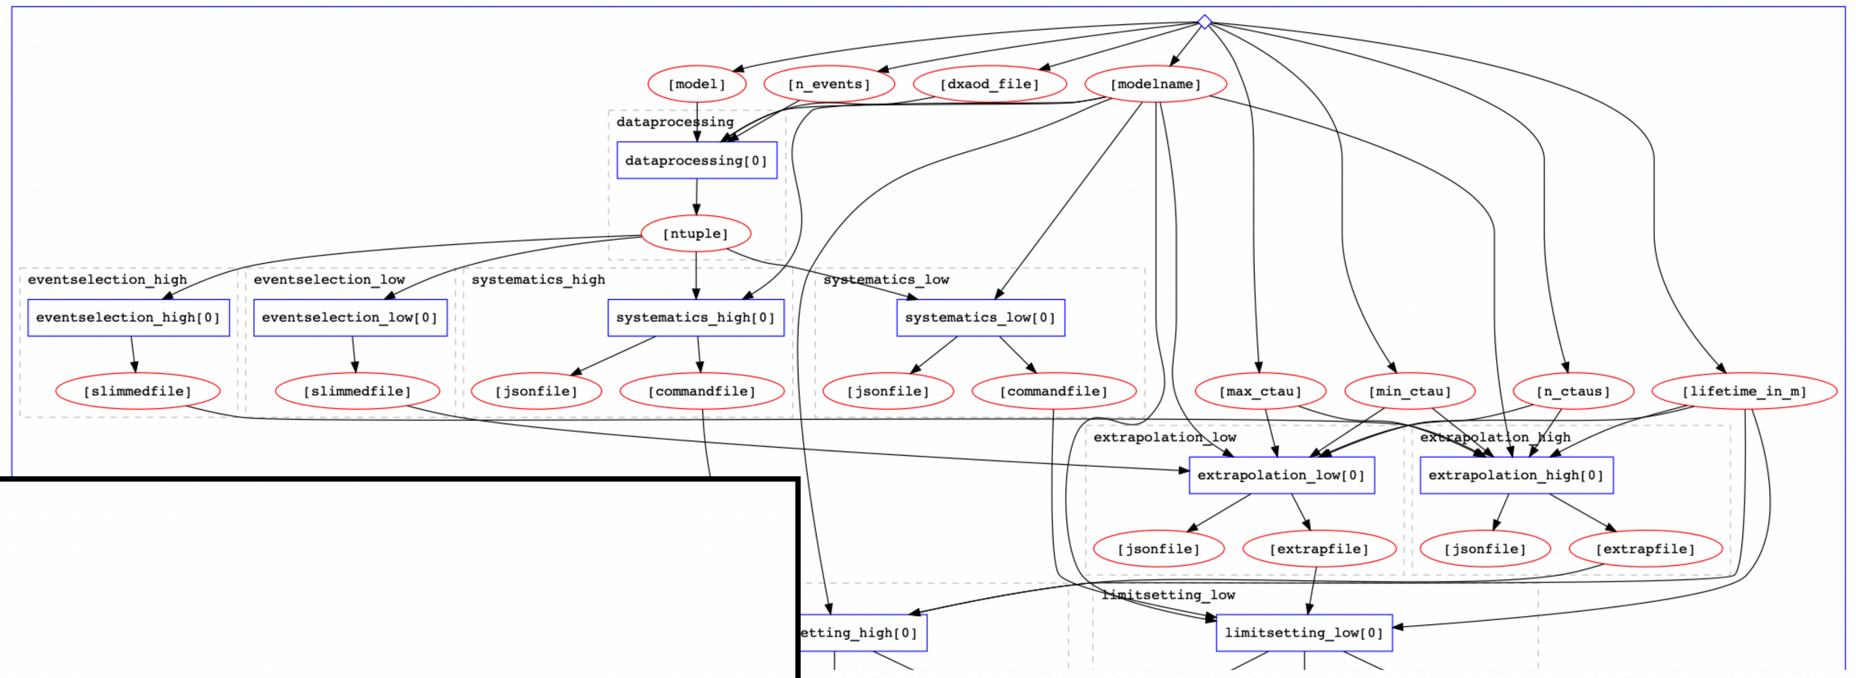

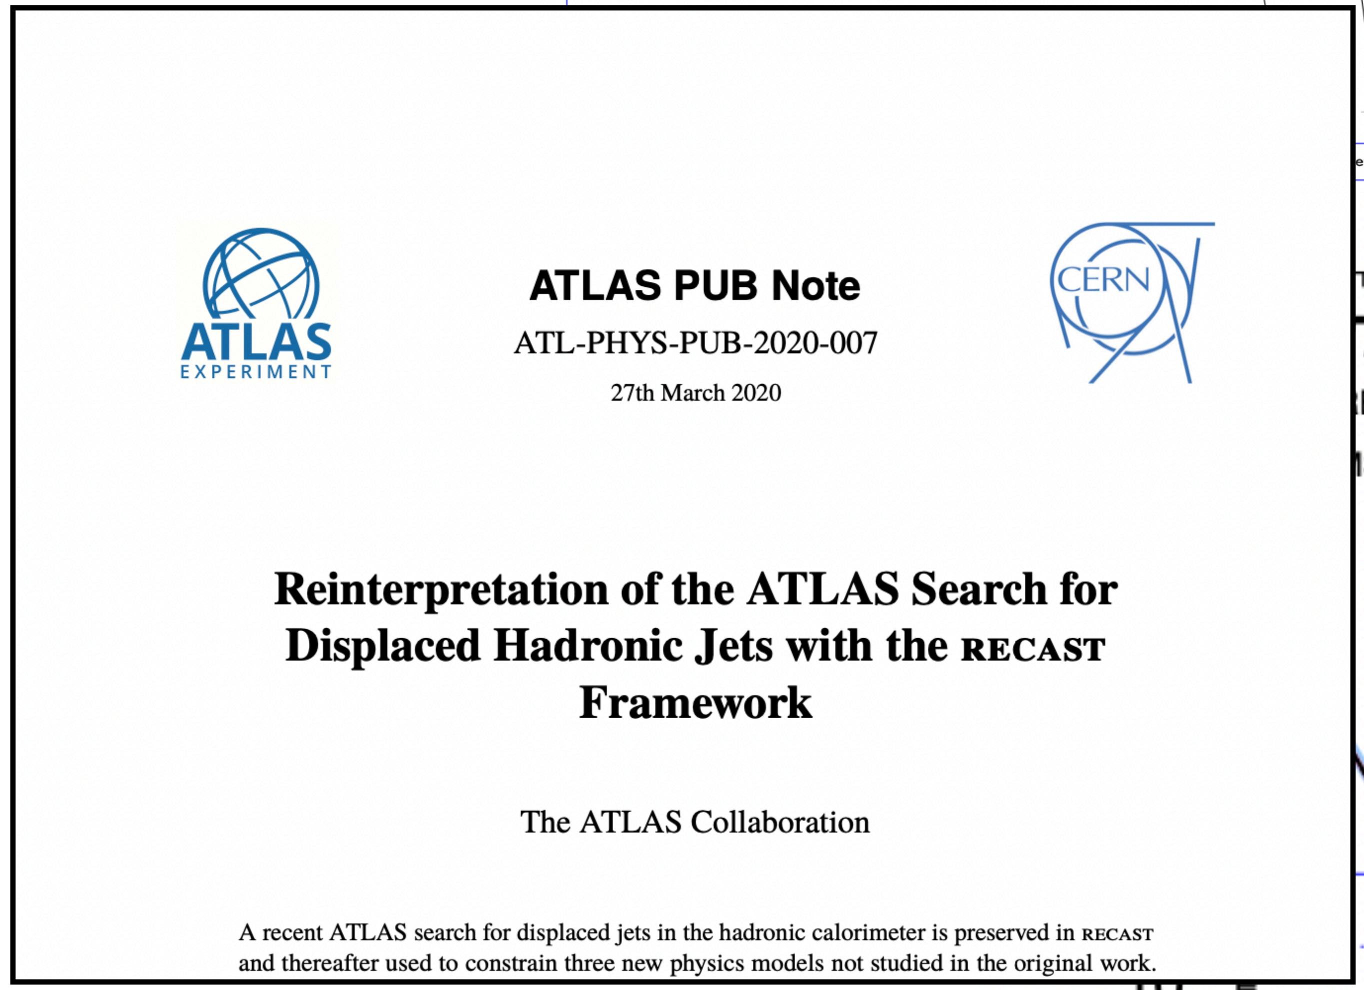

ATLAS PUB Note

ATL-PHYS-PUB-2020-007

27th March 2020

Reinterpretation of the ATLAS Search for Displaced Hadronic Jets with the RECAST Framework

The ATLAS Collaboration

A recent ATLAS search for displaced jets in the hadronic calorimeter is preserved in RECAST and thereafter used to constrain three new physics models not studied in the original work.

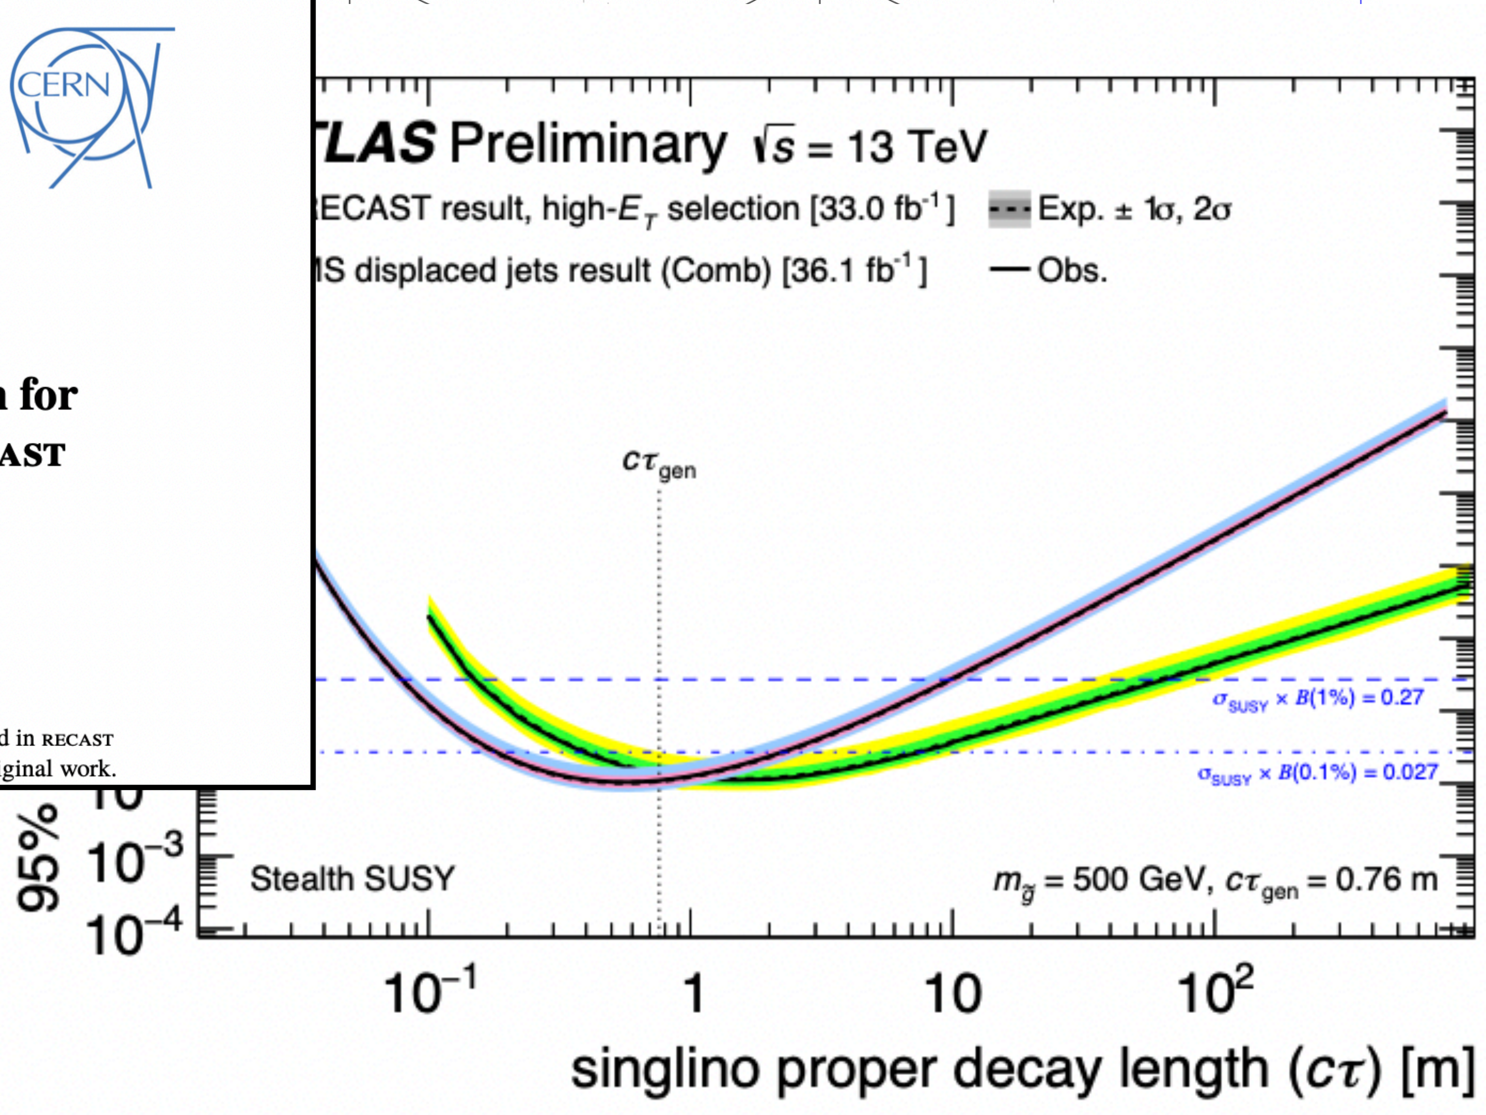

Key question: can we scale this to a general paradigm?

- e.g. study high-dimensional SUSY parameter space such as pMSSM

Services as Kubernetes Applications

From a Kubernetes point-of-view complex services like REANA are just another "application": easy to deploy

```
helm repo add reanahub https://reanahub.github.io/reana
helm install --devel reana reanahub/reana --wait
```

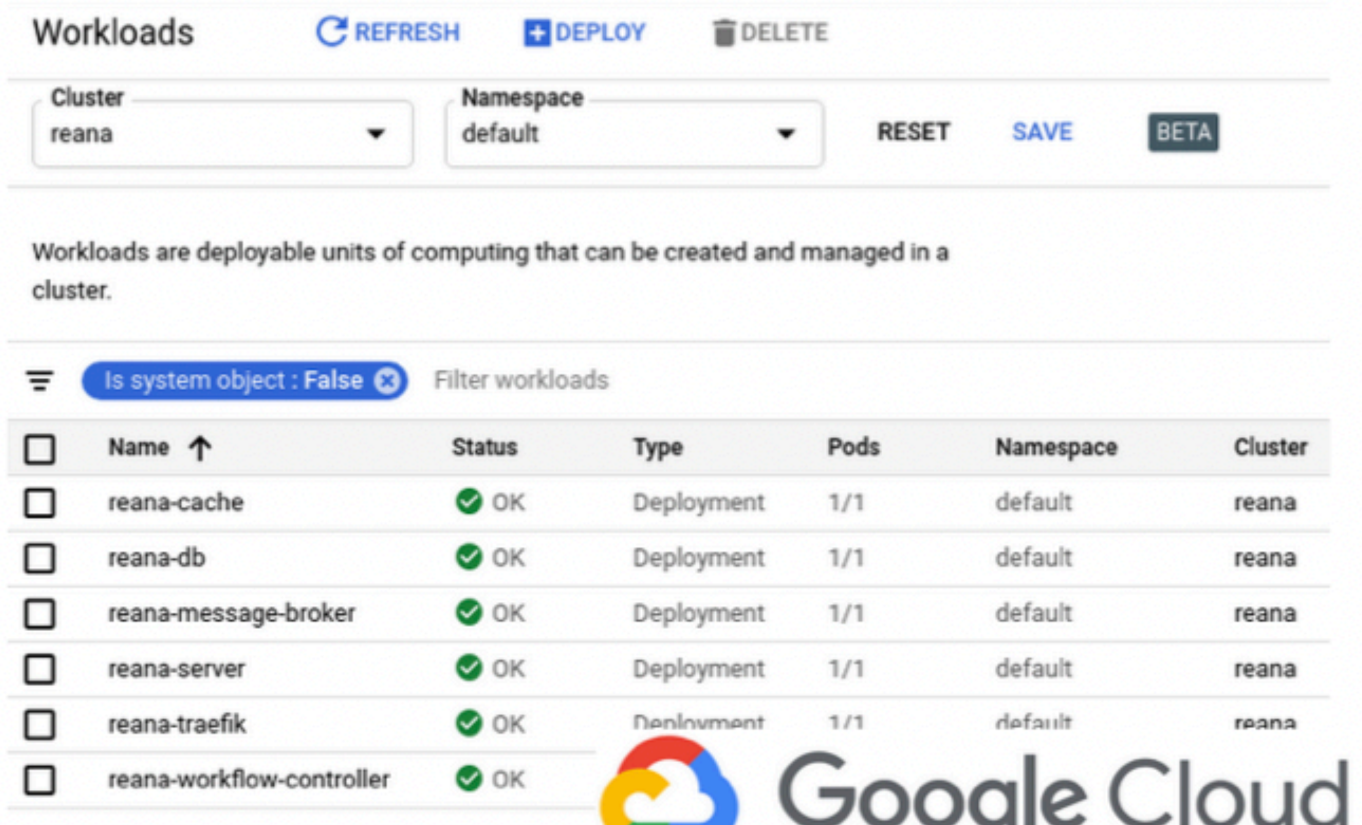

REANA on Google Cloud

Same installation at CERN
and Google Cloud thanks to uniformity
of Kubernetes APIs

Scaling REANA

To study 19-D pMSSM one would need to run O(100k) analysis runs

- processing many samples, complex fits.. great stress test for REANA
- scaled at CERN to O(1k cores) looking to push further (a la grid)
- develop basis for production use of REANA inside of collaborations

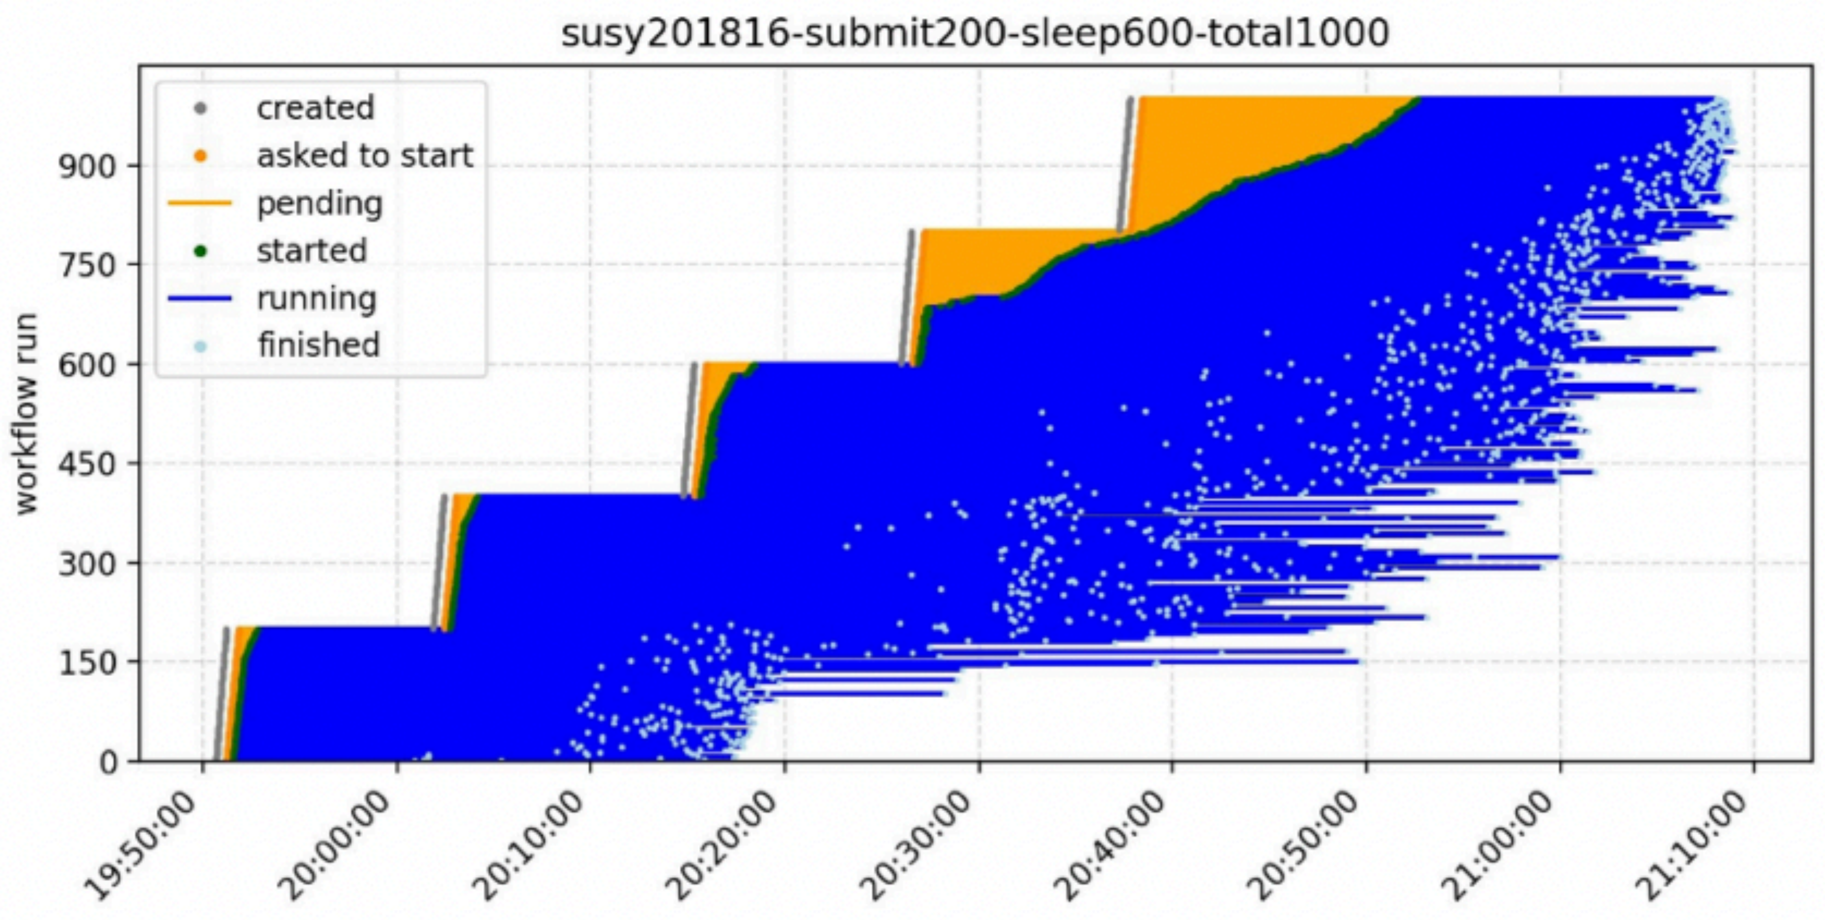

Old cluster (448 cores)

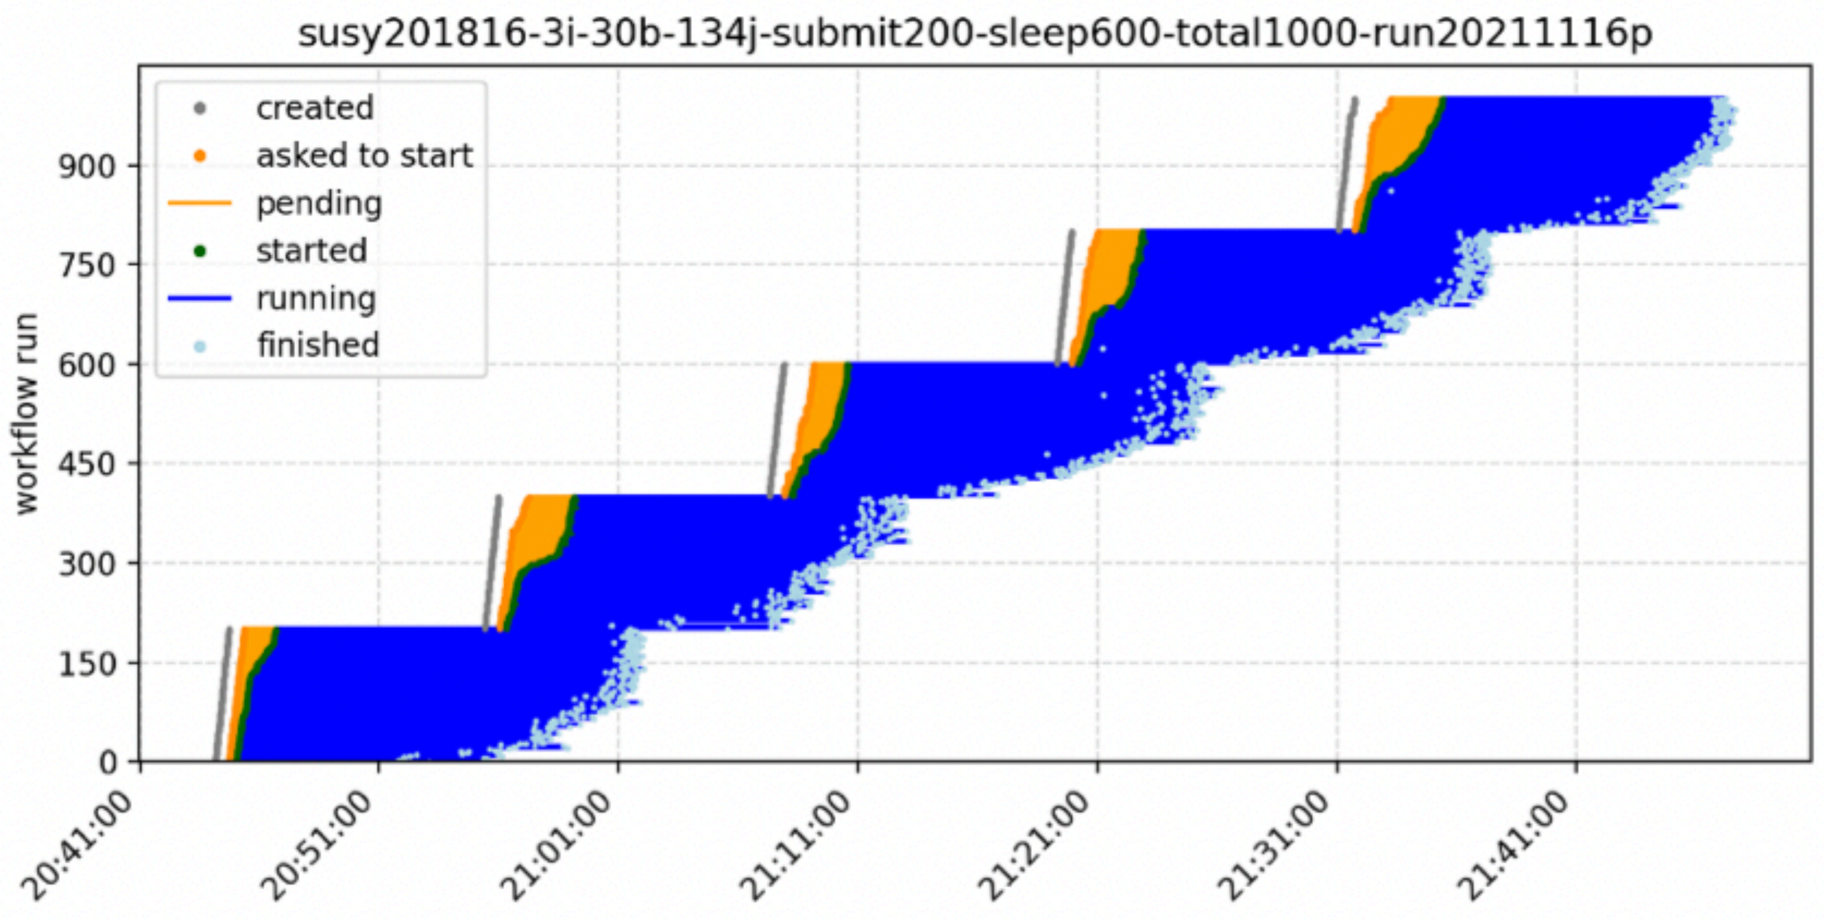

New cluster (1072 cores)

Beyond Simple Batch

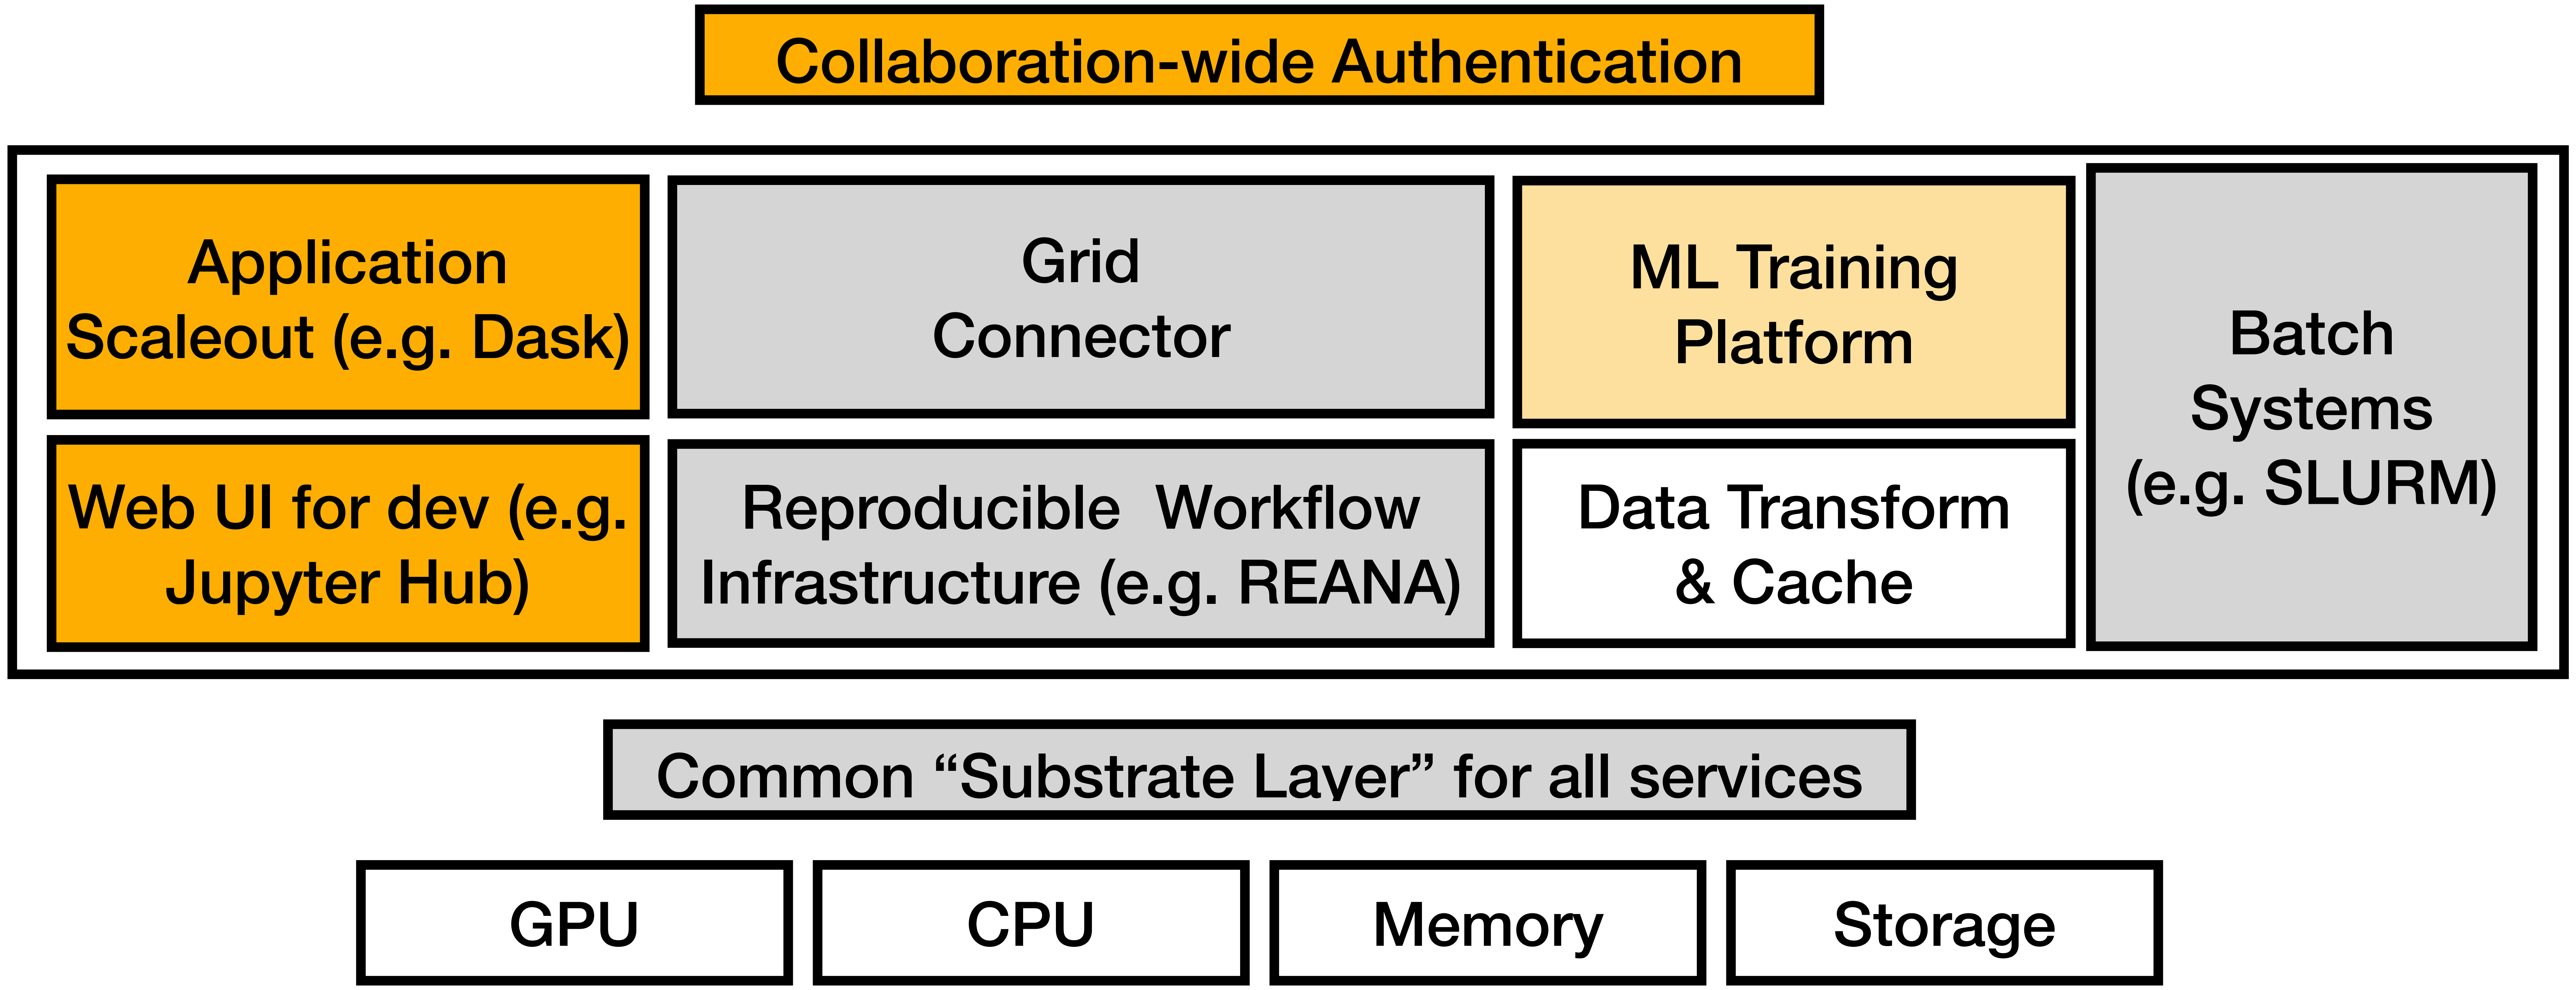

Interactive Analysis

So far, everything was batch (either direct submission or workflows)

But to analysis should be explorative, interactive & fast turnaround at HL-LHC scales. Blast from the Past:

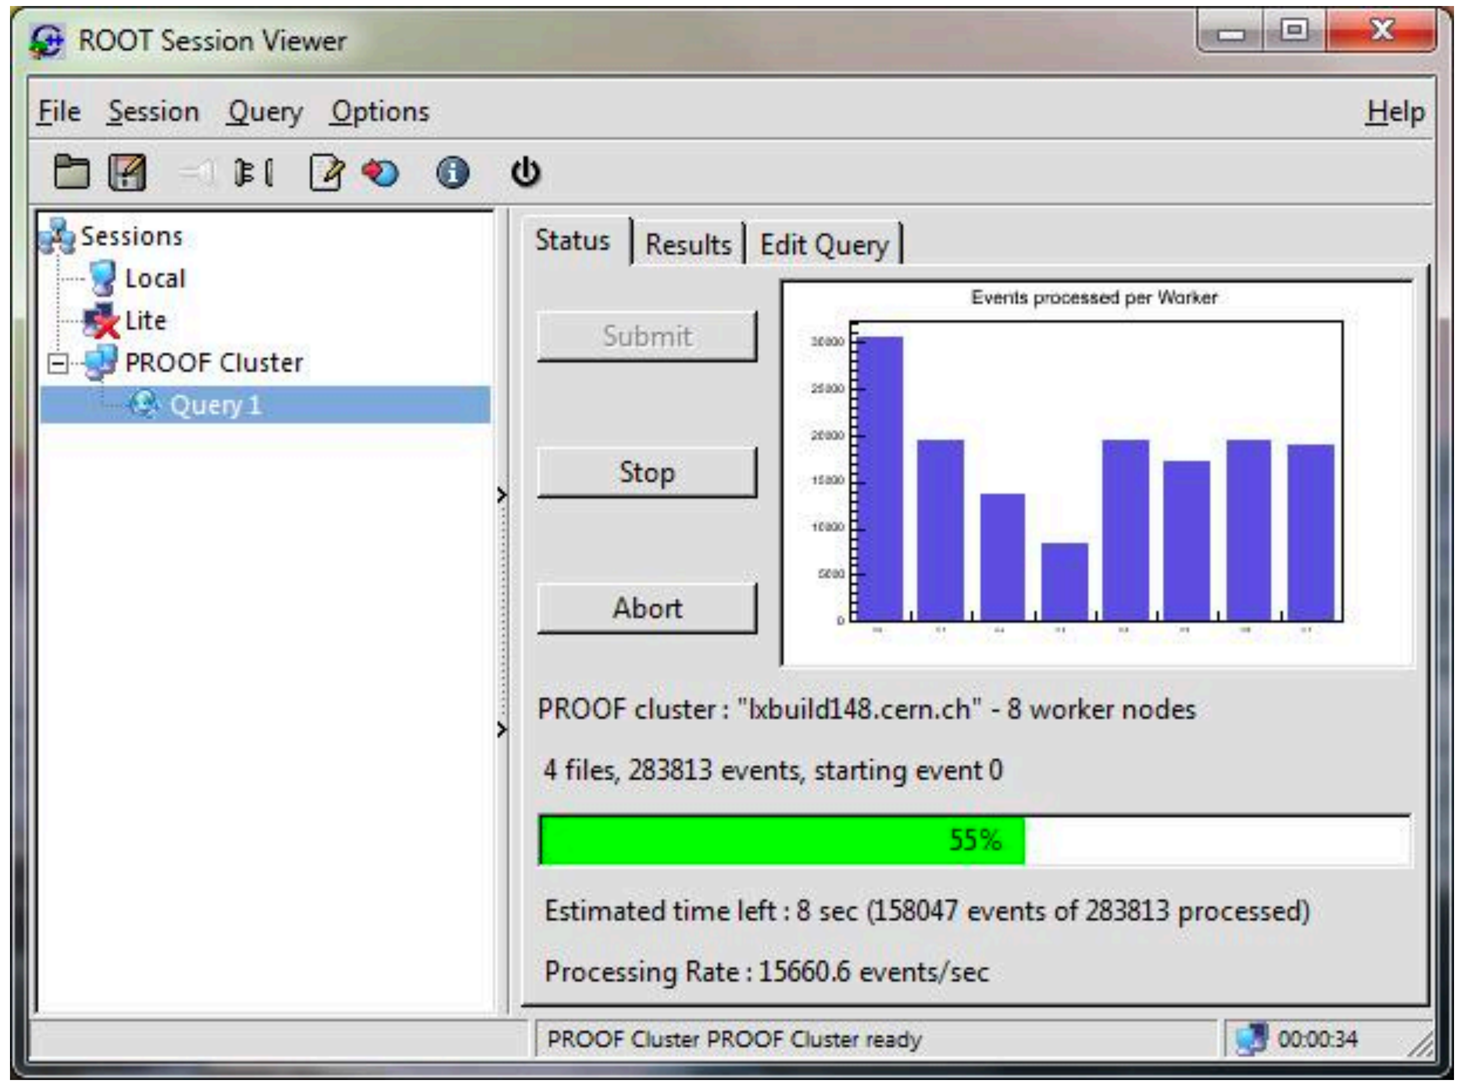

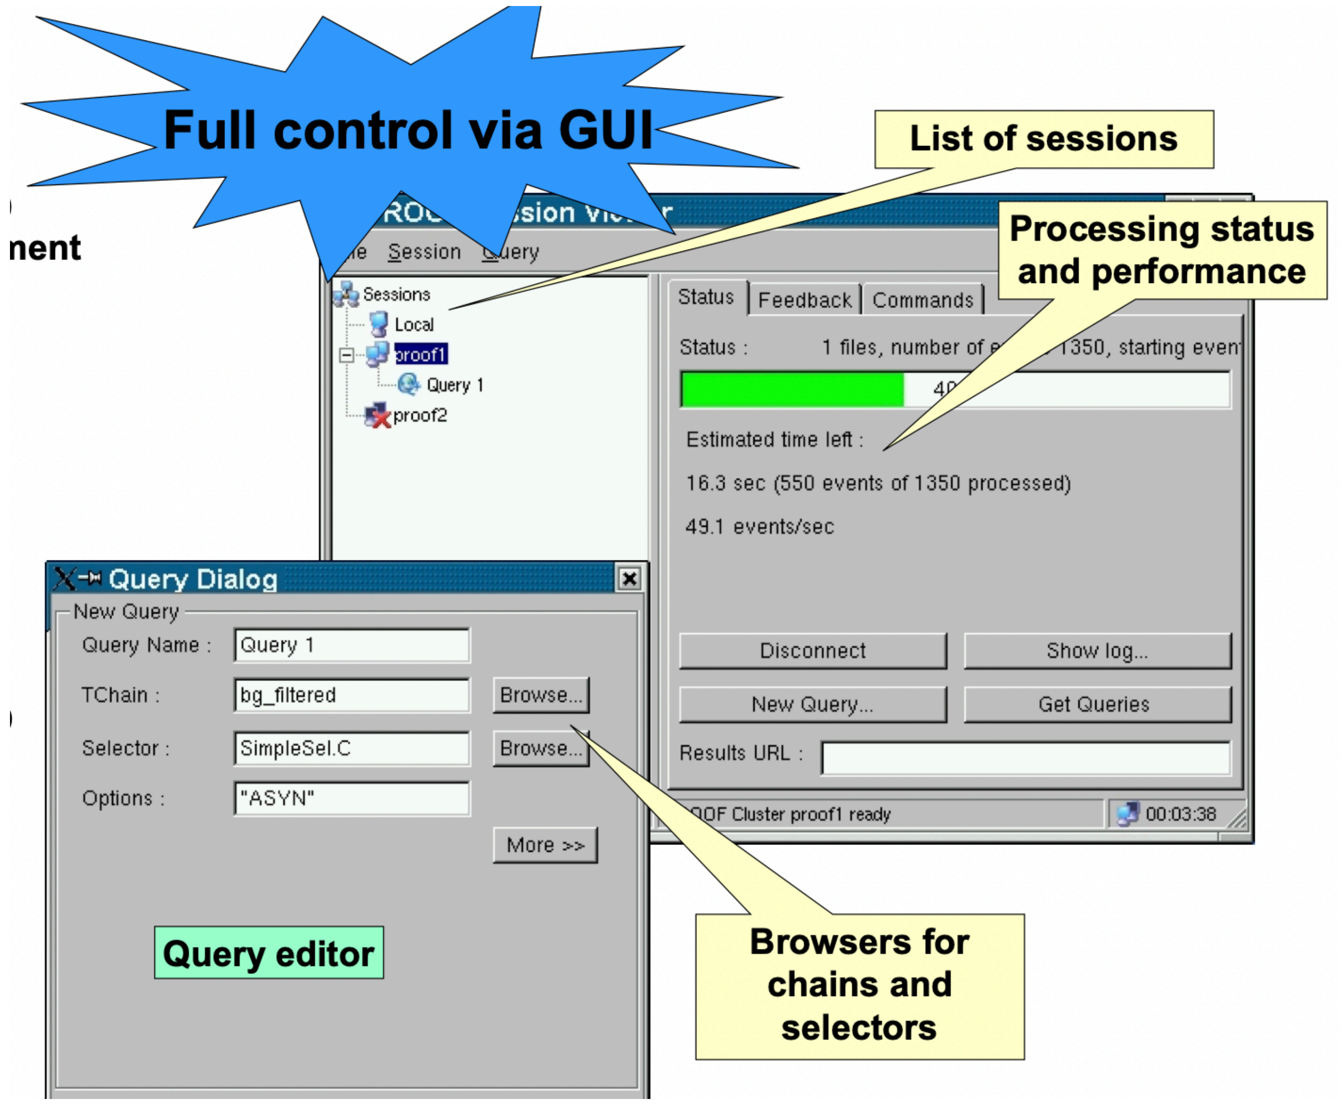

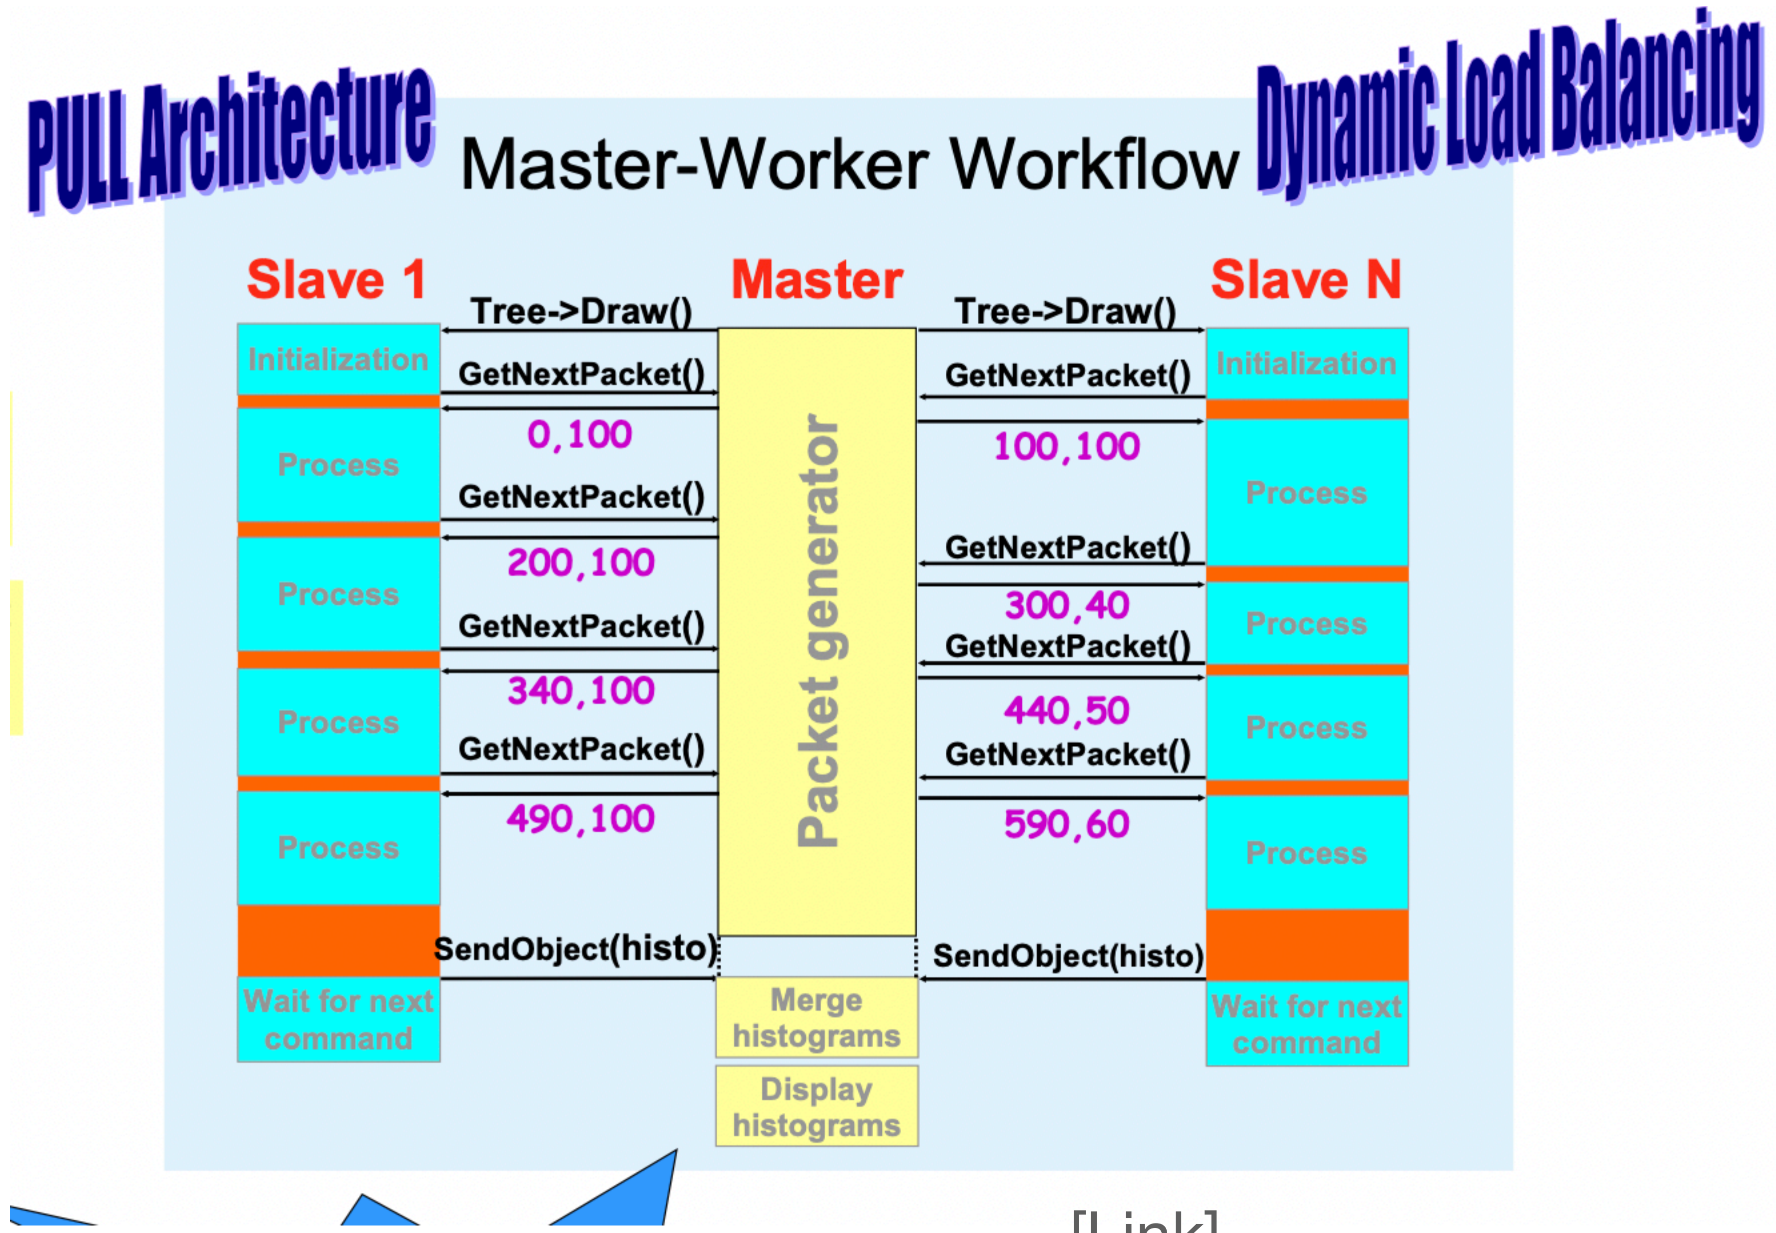

[Link]

PROOF had the right vision, but technology wasn't really there

A modern take on interactive analysis

Data Science community is investing a lot in “distributed data frame” processors for large-scale data analysis. Good example: Dask

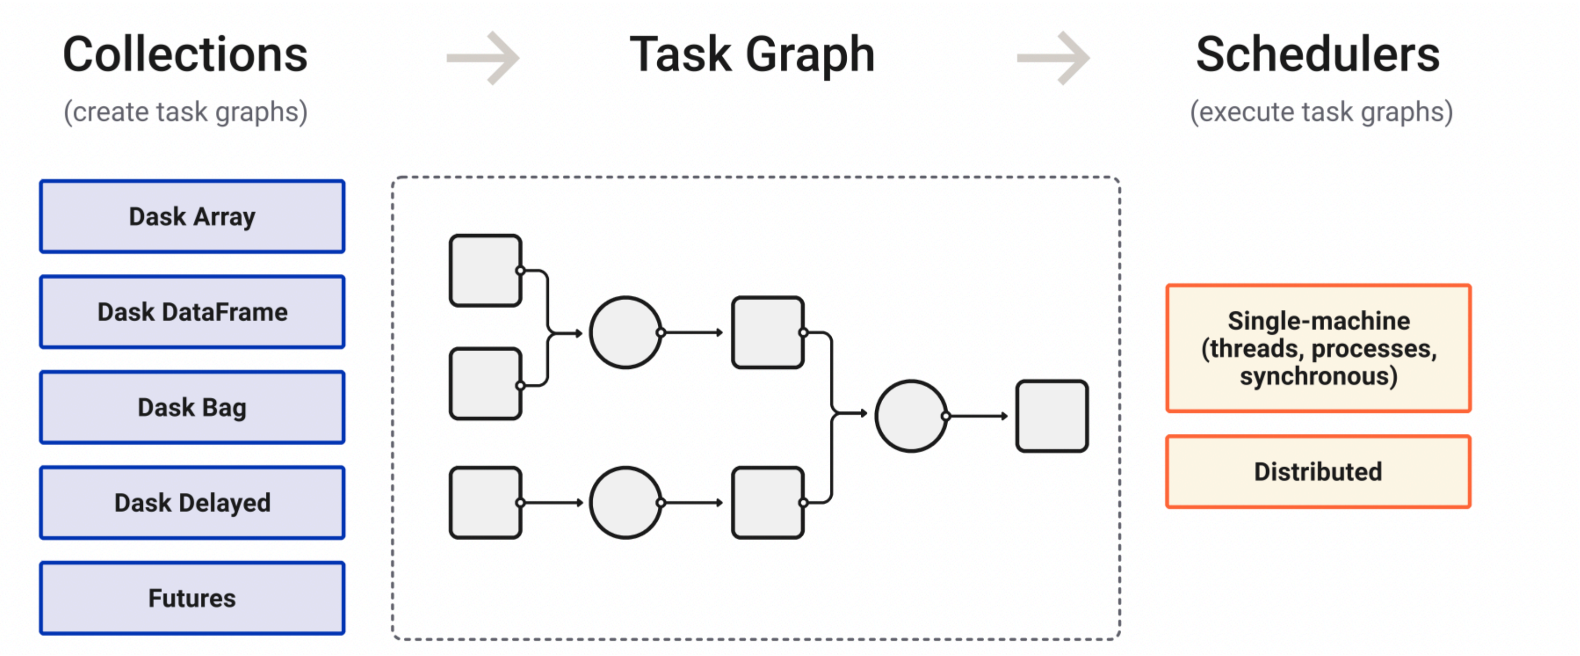

DASK

Web-based technologies instead of Native GUIs: Jupyter, VS Code

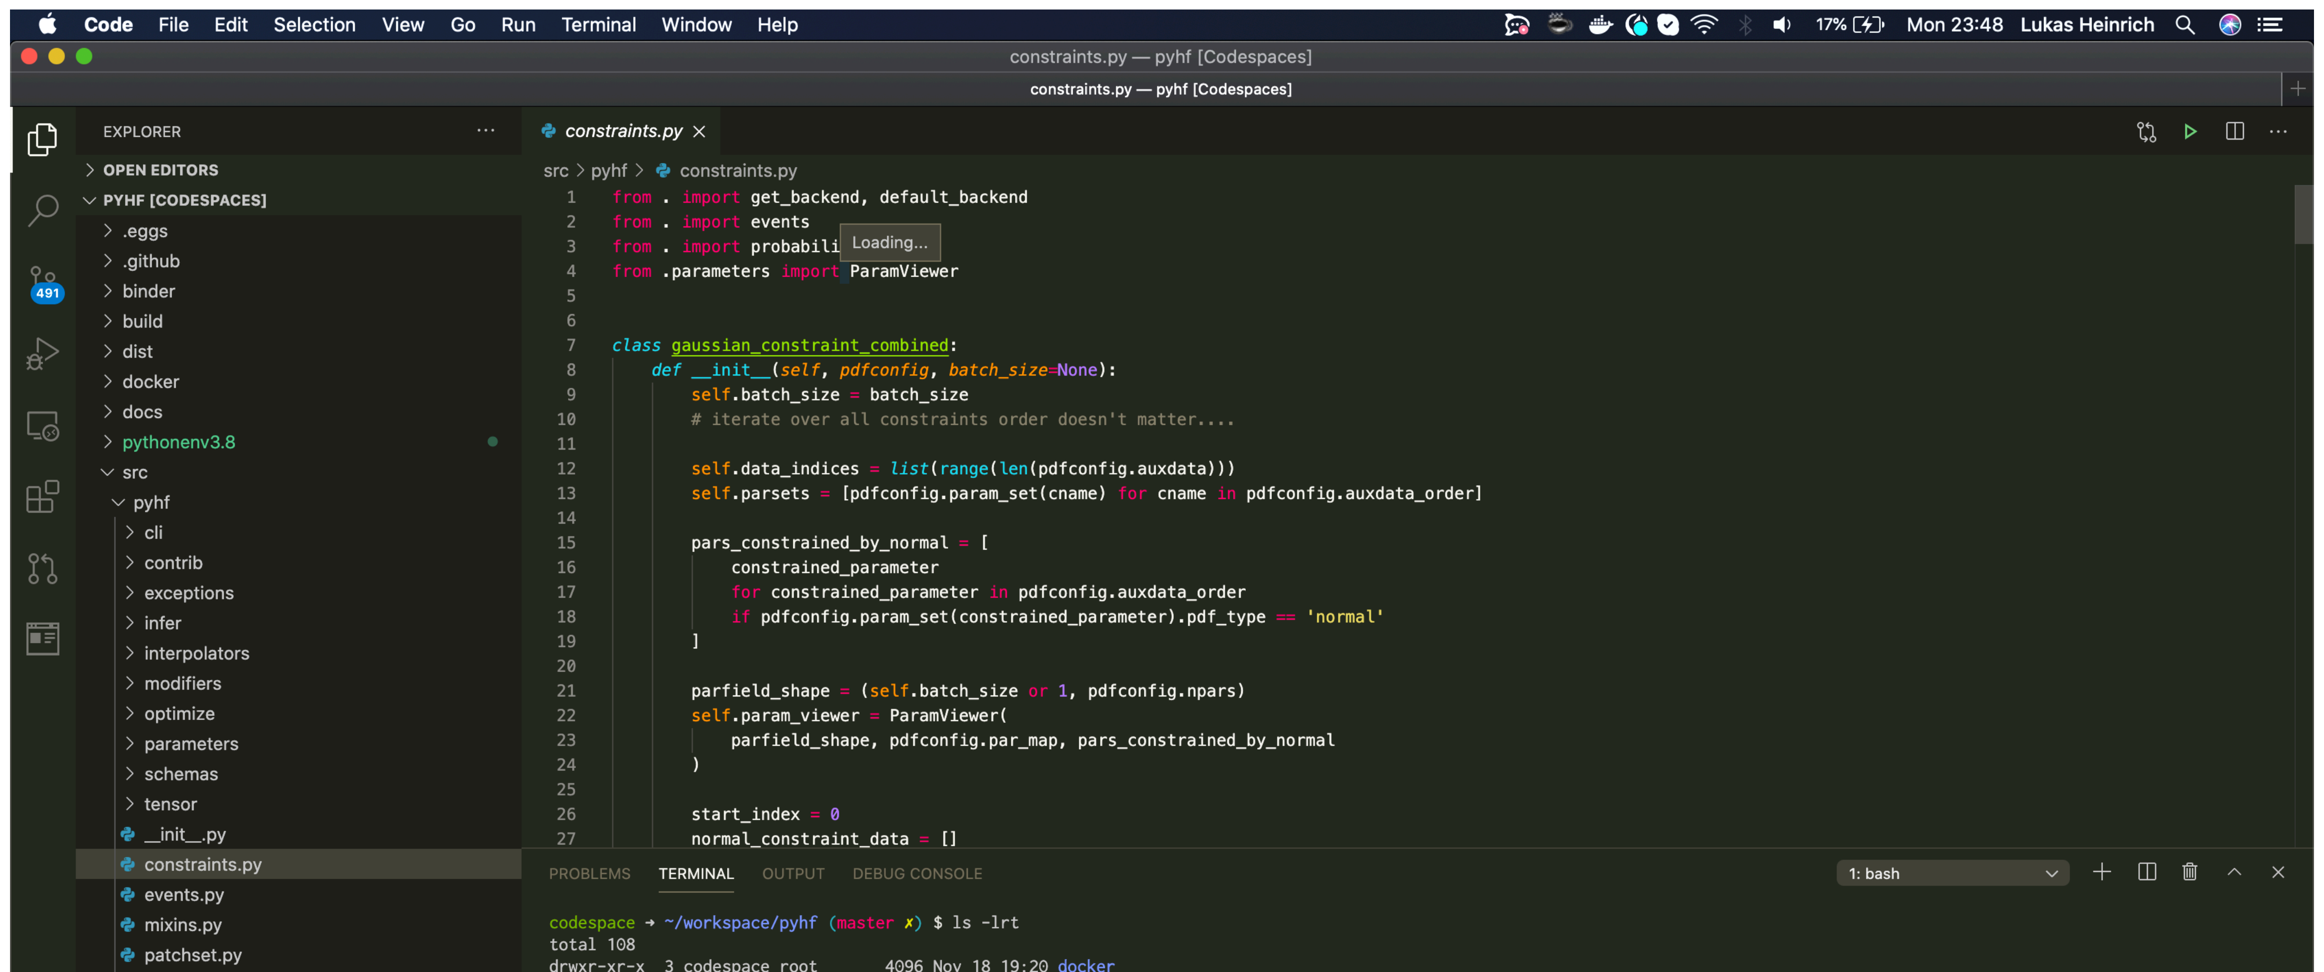

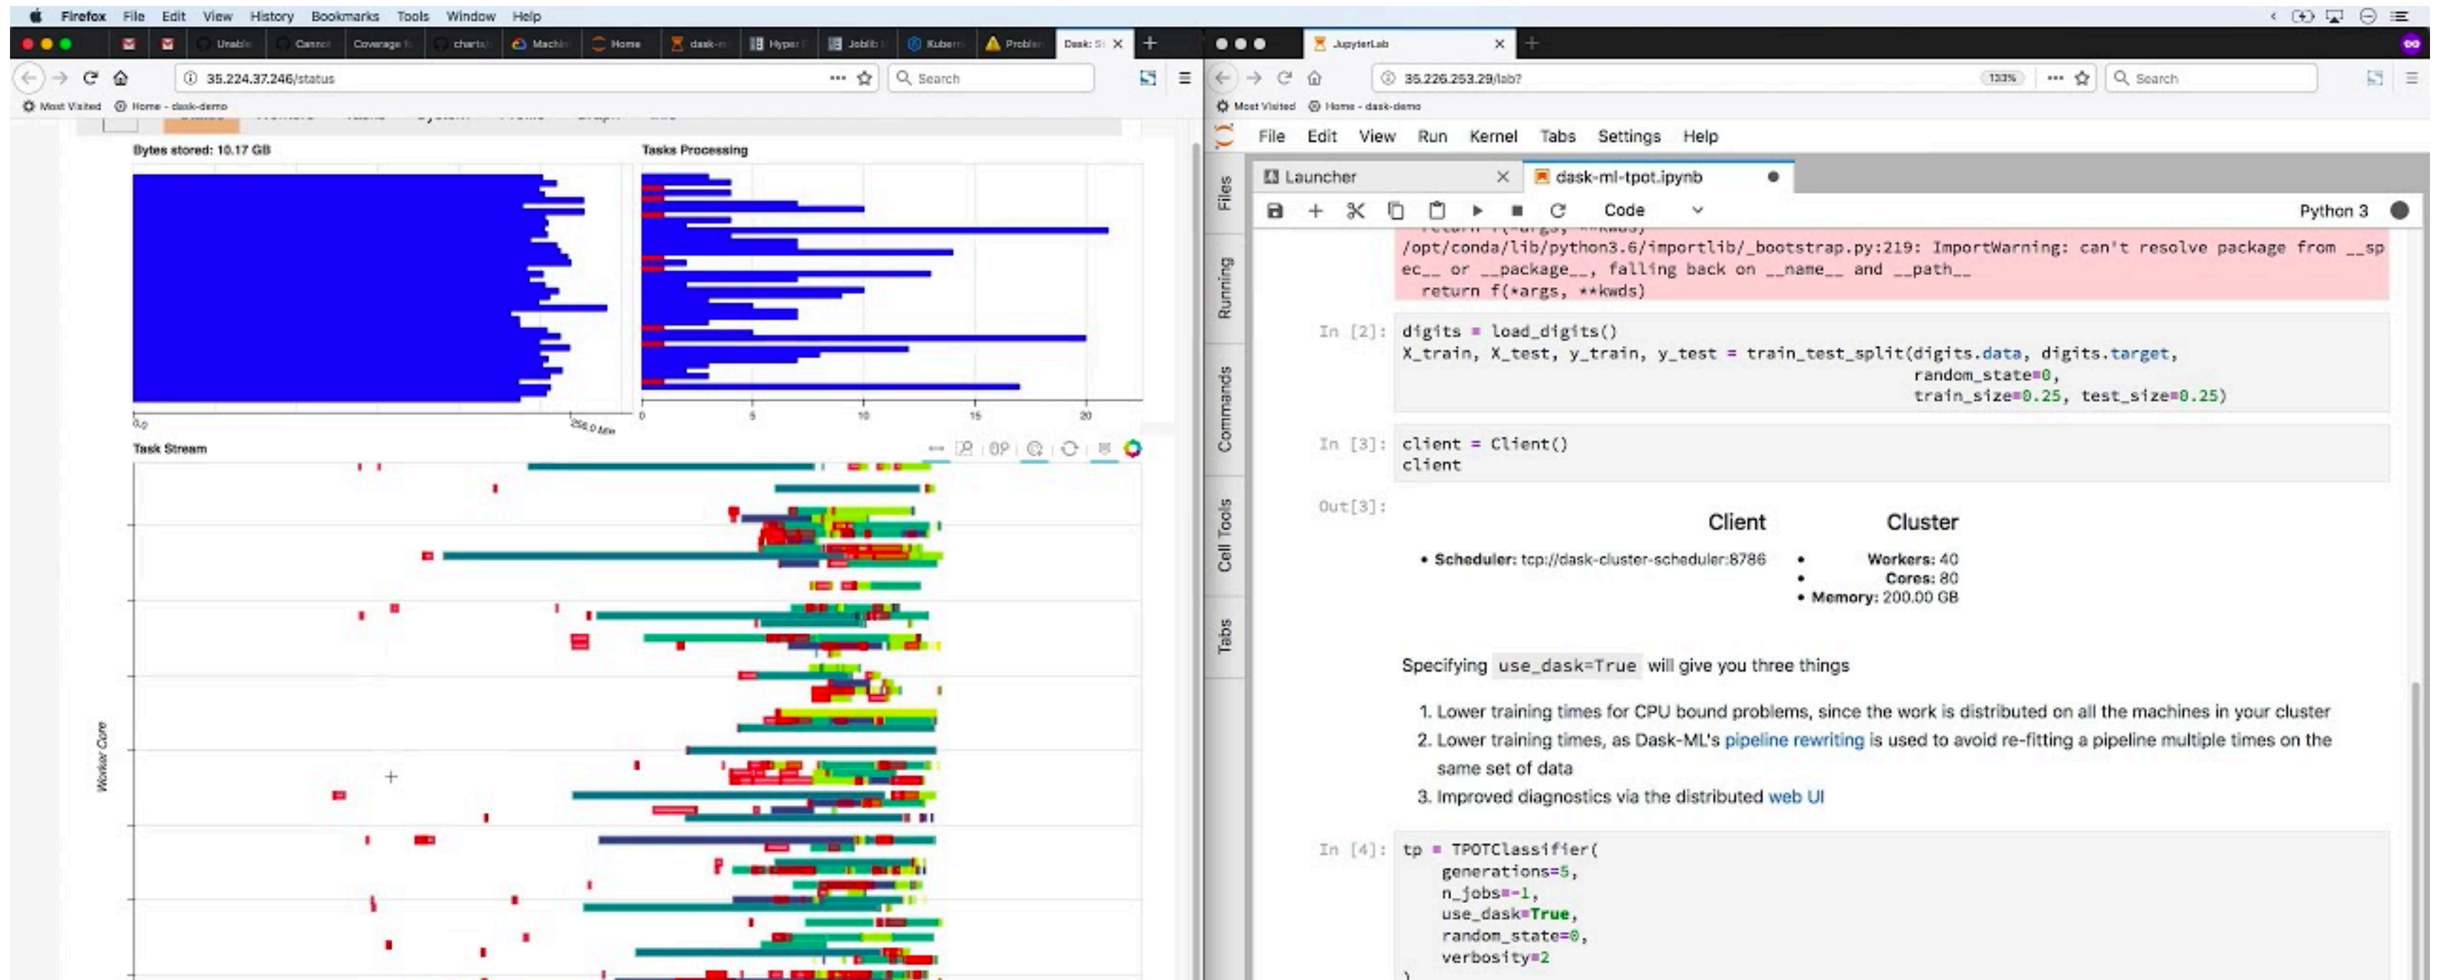

Interactive Analysis for HEP

Closely linked to developments for pythonic HEP analysis

- easy to do “demos”, but for realistic analysis require a full solution for systematic uncertainties, …

IRIS-HEP Analysis Grand Challenge: demonstrate large-scale and realistic data analysis on Open Data (the “A” in FAIR”)

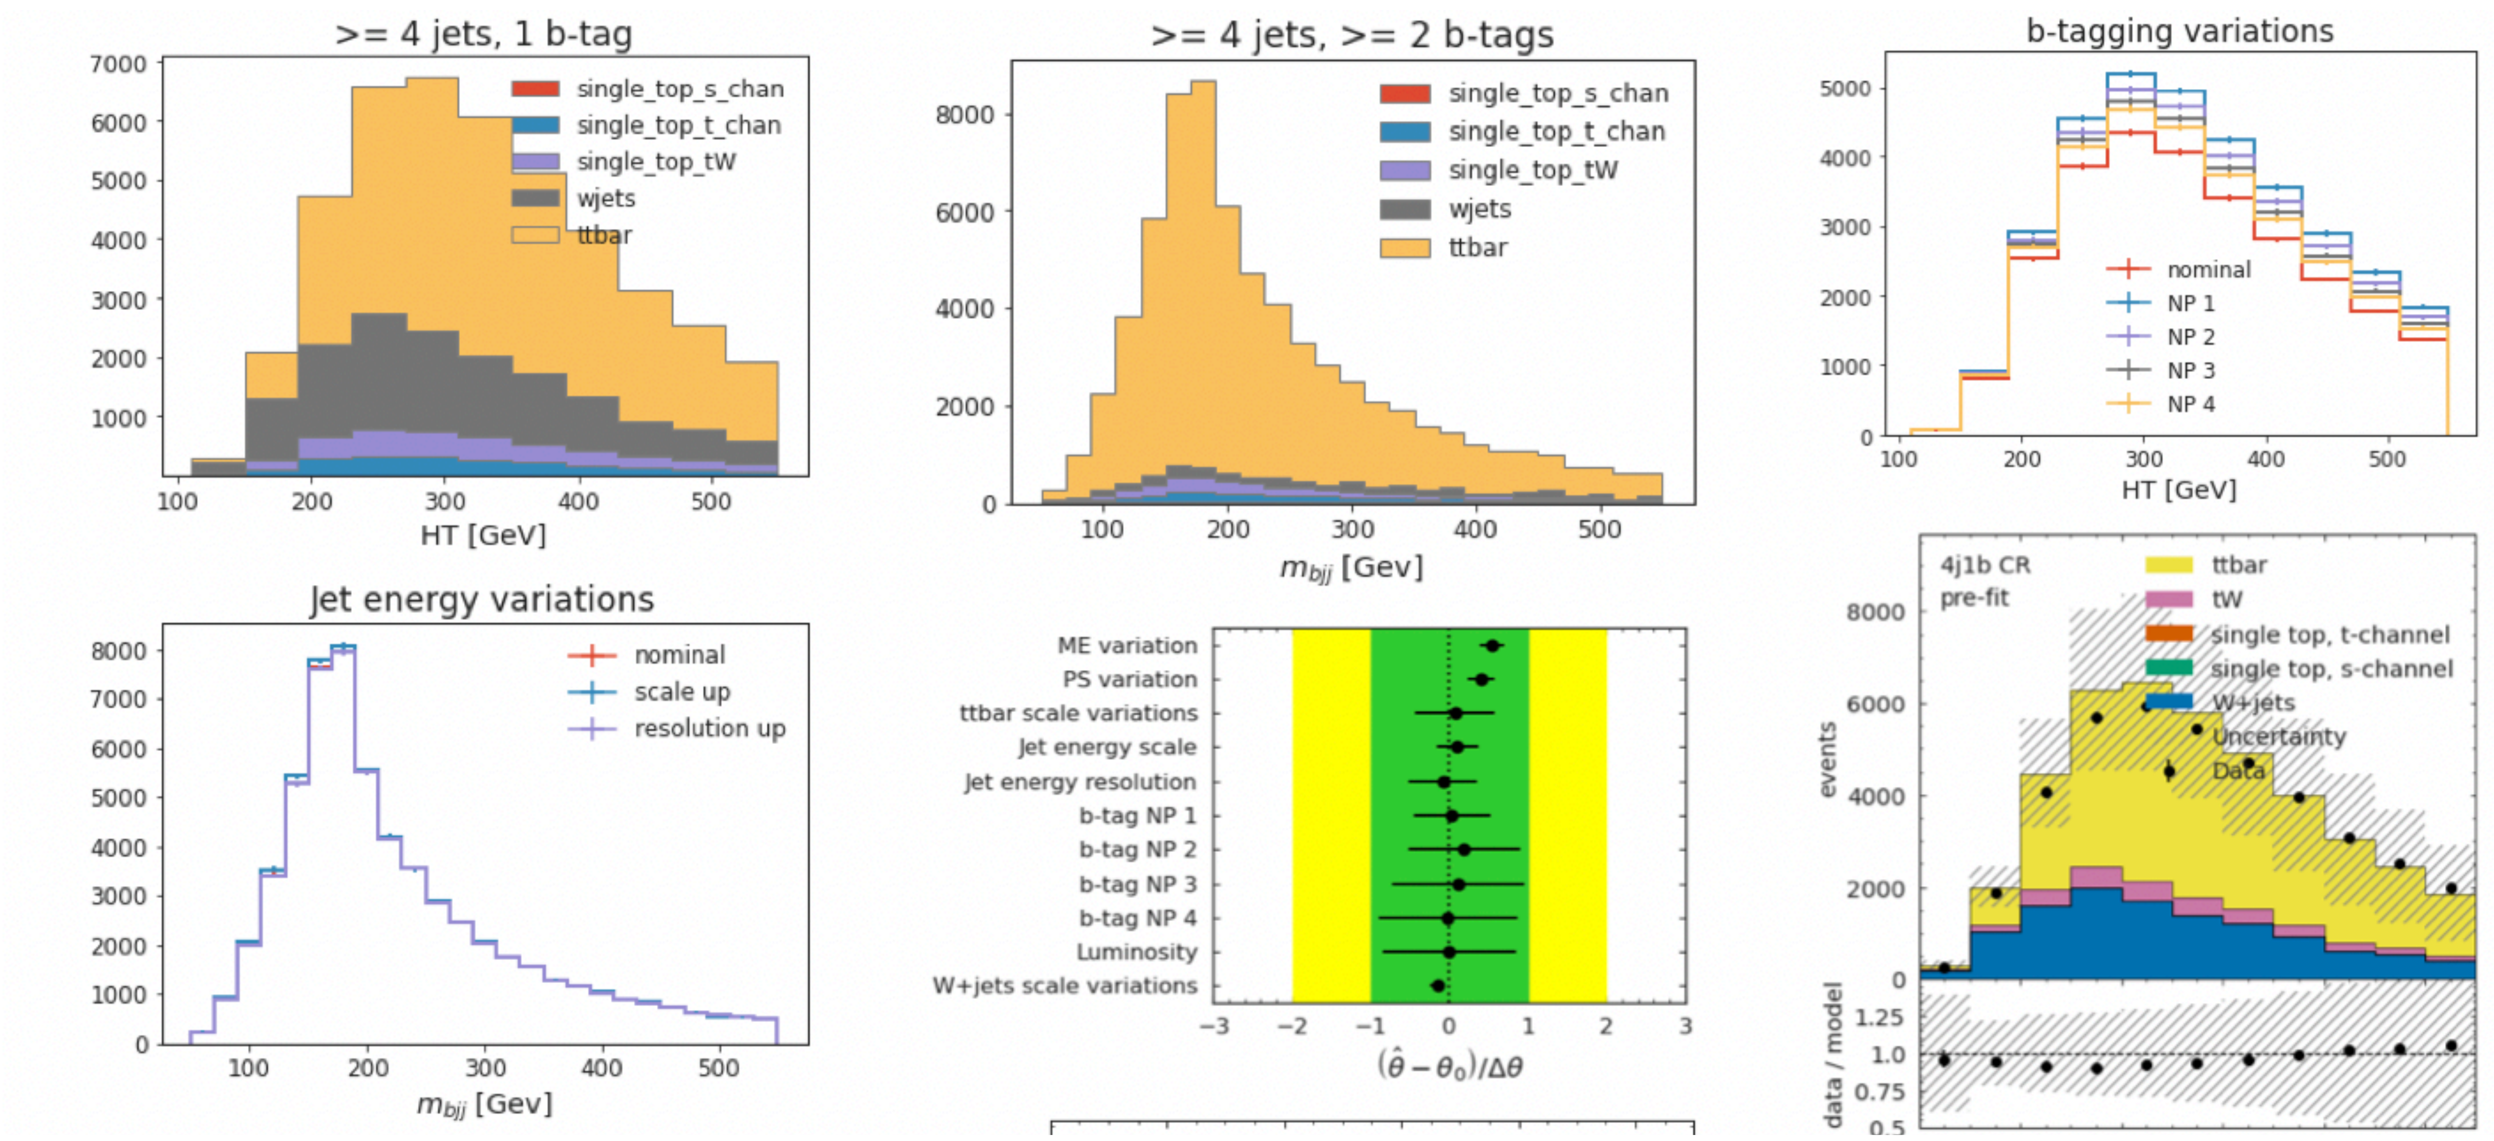

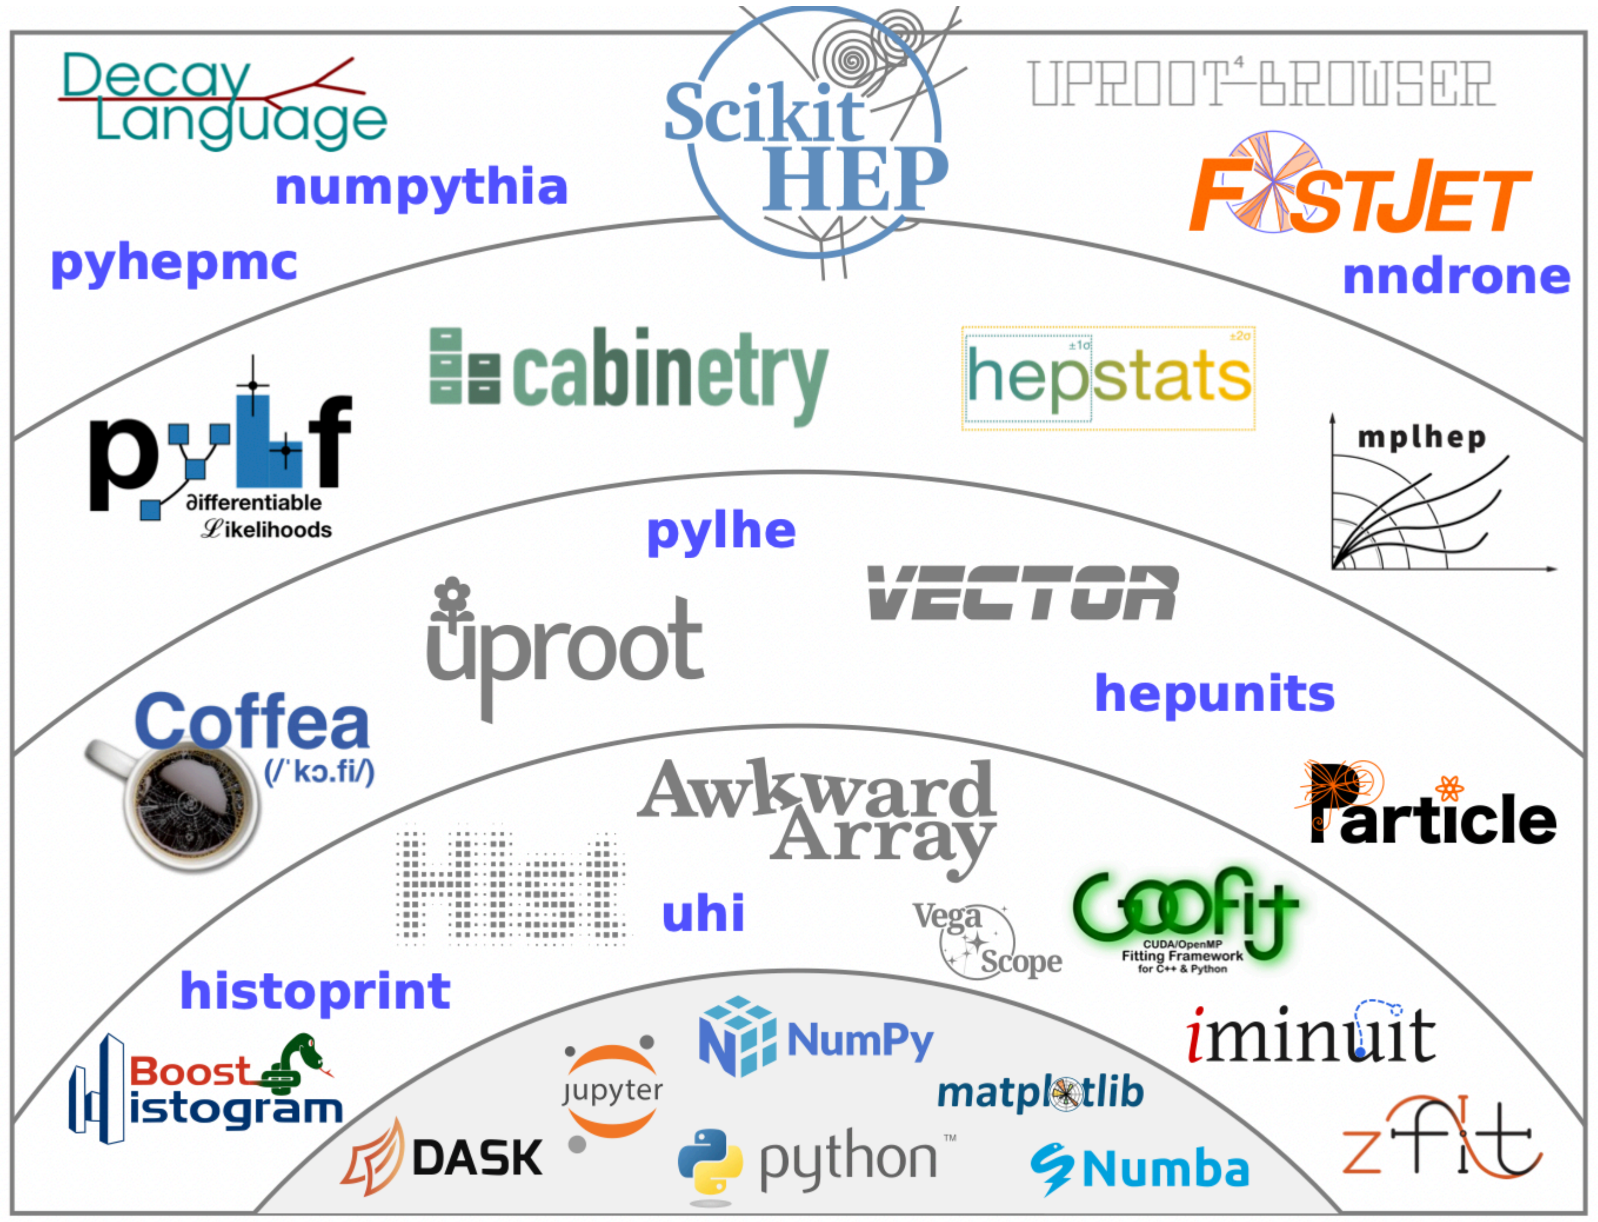

A recurring Picture

Again, “Jupyter” and “Dask” are just add. applications for Kubernetes like REANA and Harvester: Install via Helm & Scale dynamically

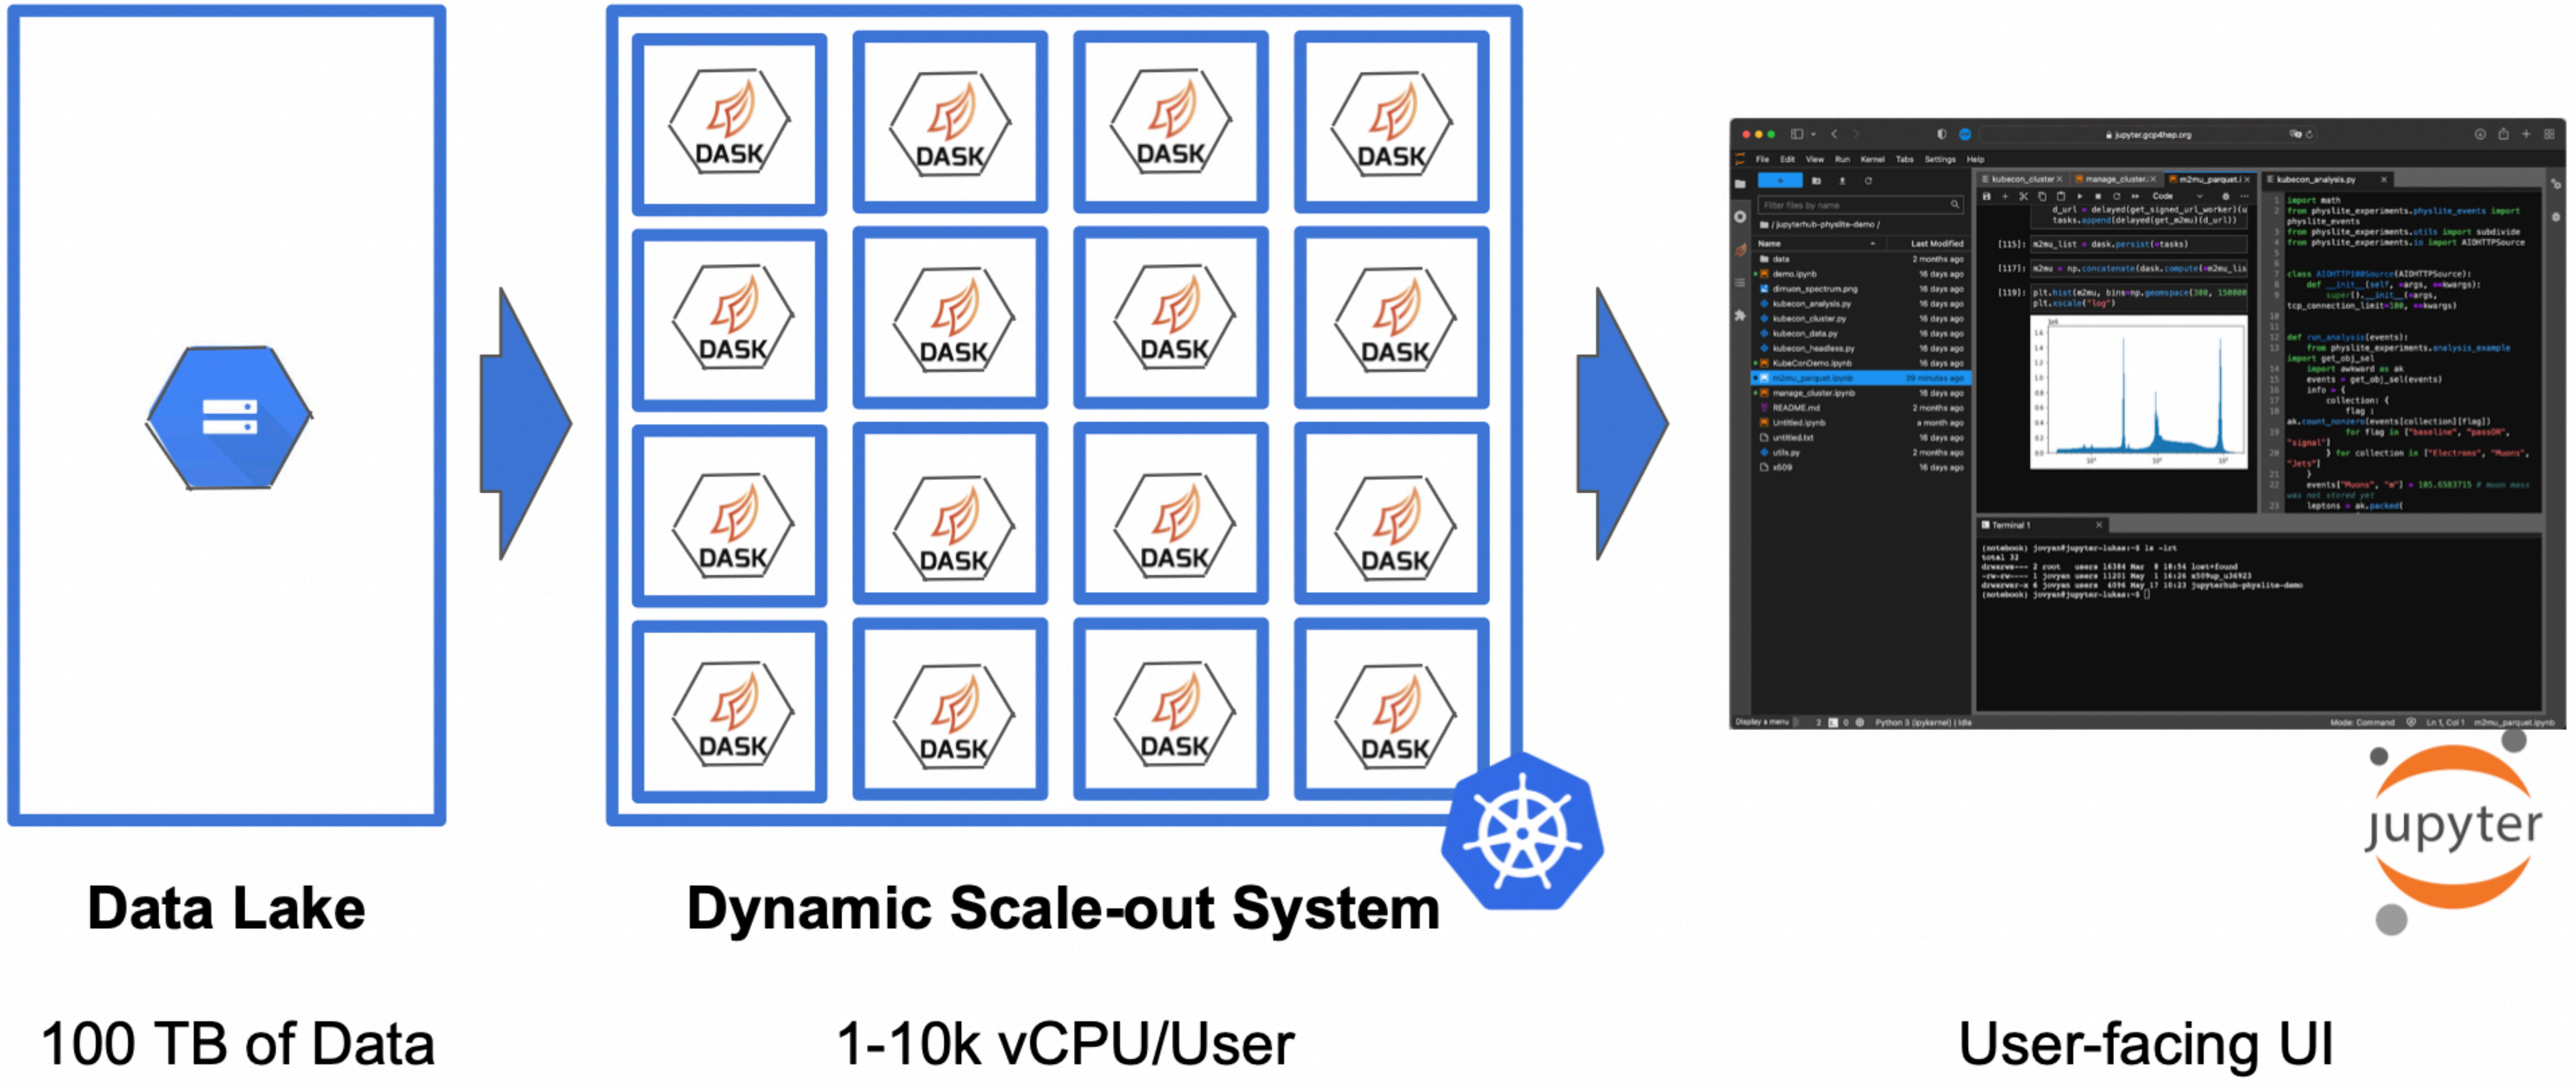

Scale on Google Cloud

Deploy Jupyter + Dask with VO-wide authentication.

- All ATLAS users could get access
- Data via Rucio-on-Google (to minimize data transfer in/out of cloud)
- Custom environments for ML (GPU + ML frameworks)

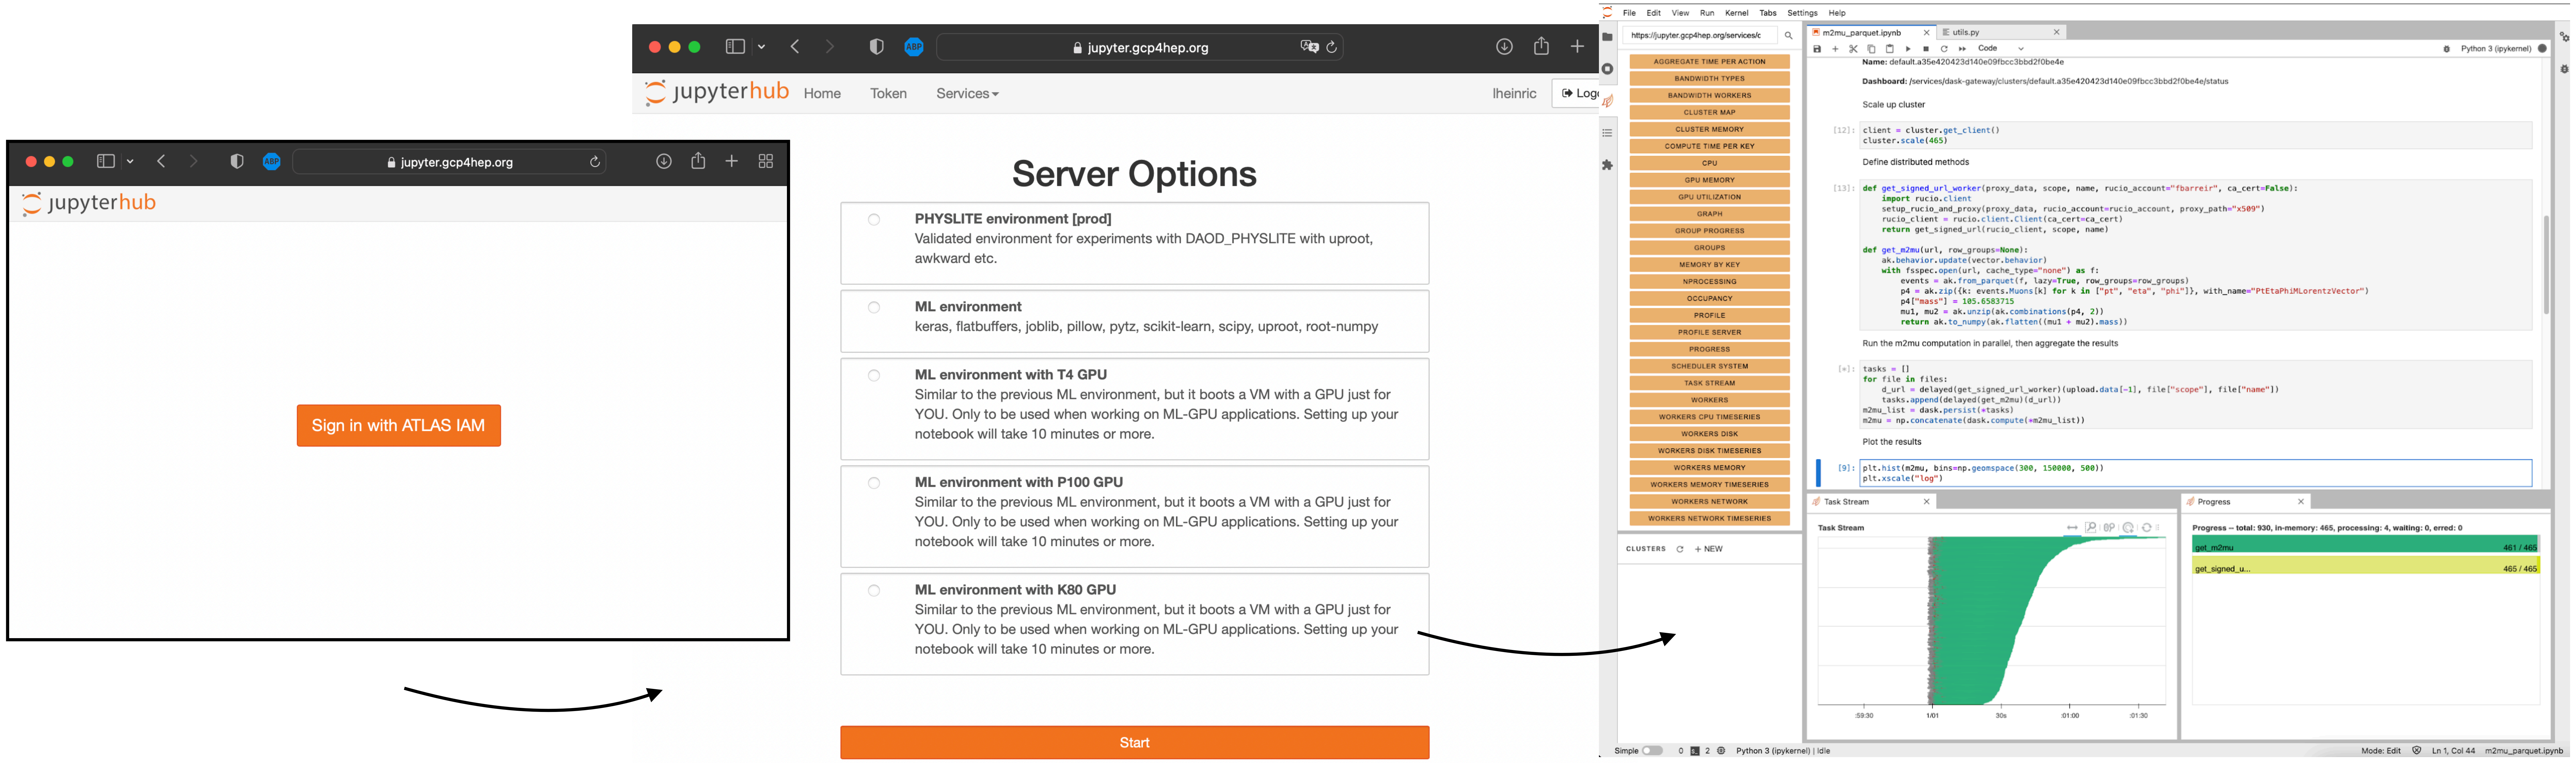

A modern Take

Exploit elasticity of cloud: scale for O(1000) cores for a user analysis
Example: Throughput test for μμ-invariant mass of full Run-2
on ATLAS data in minutes

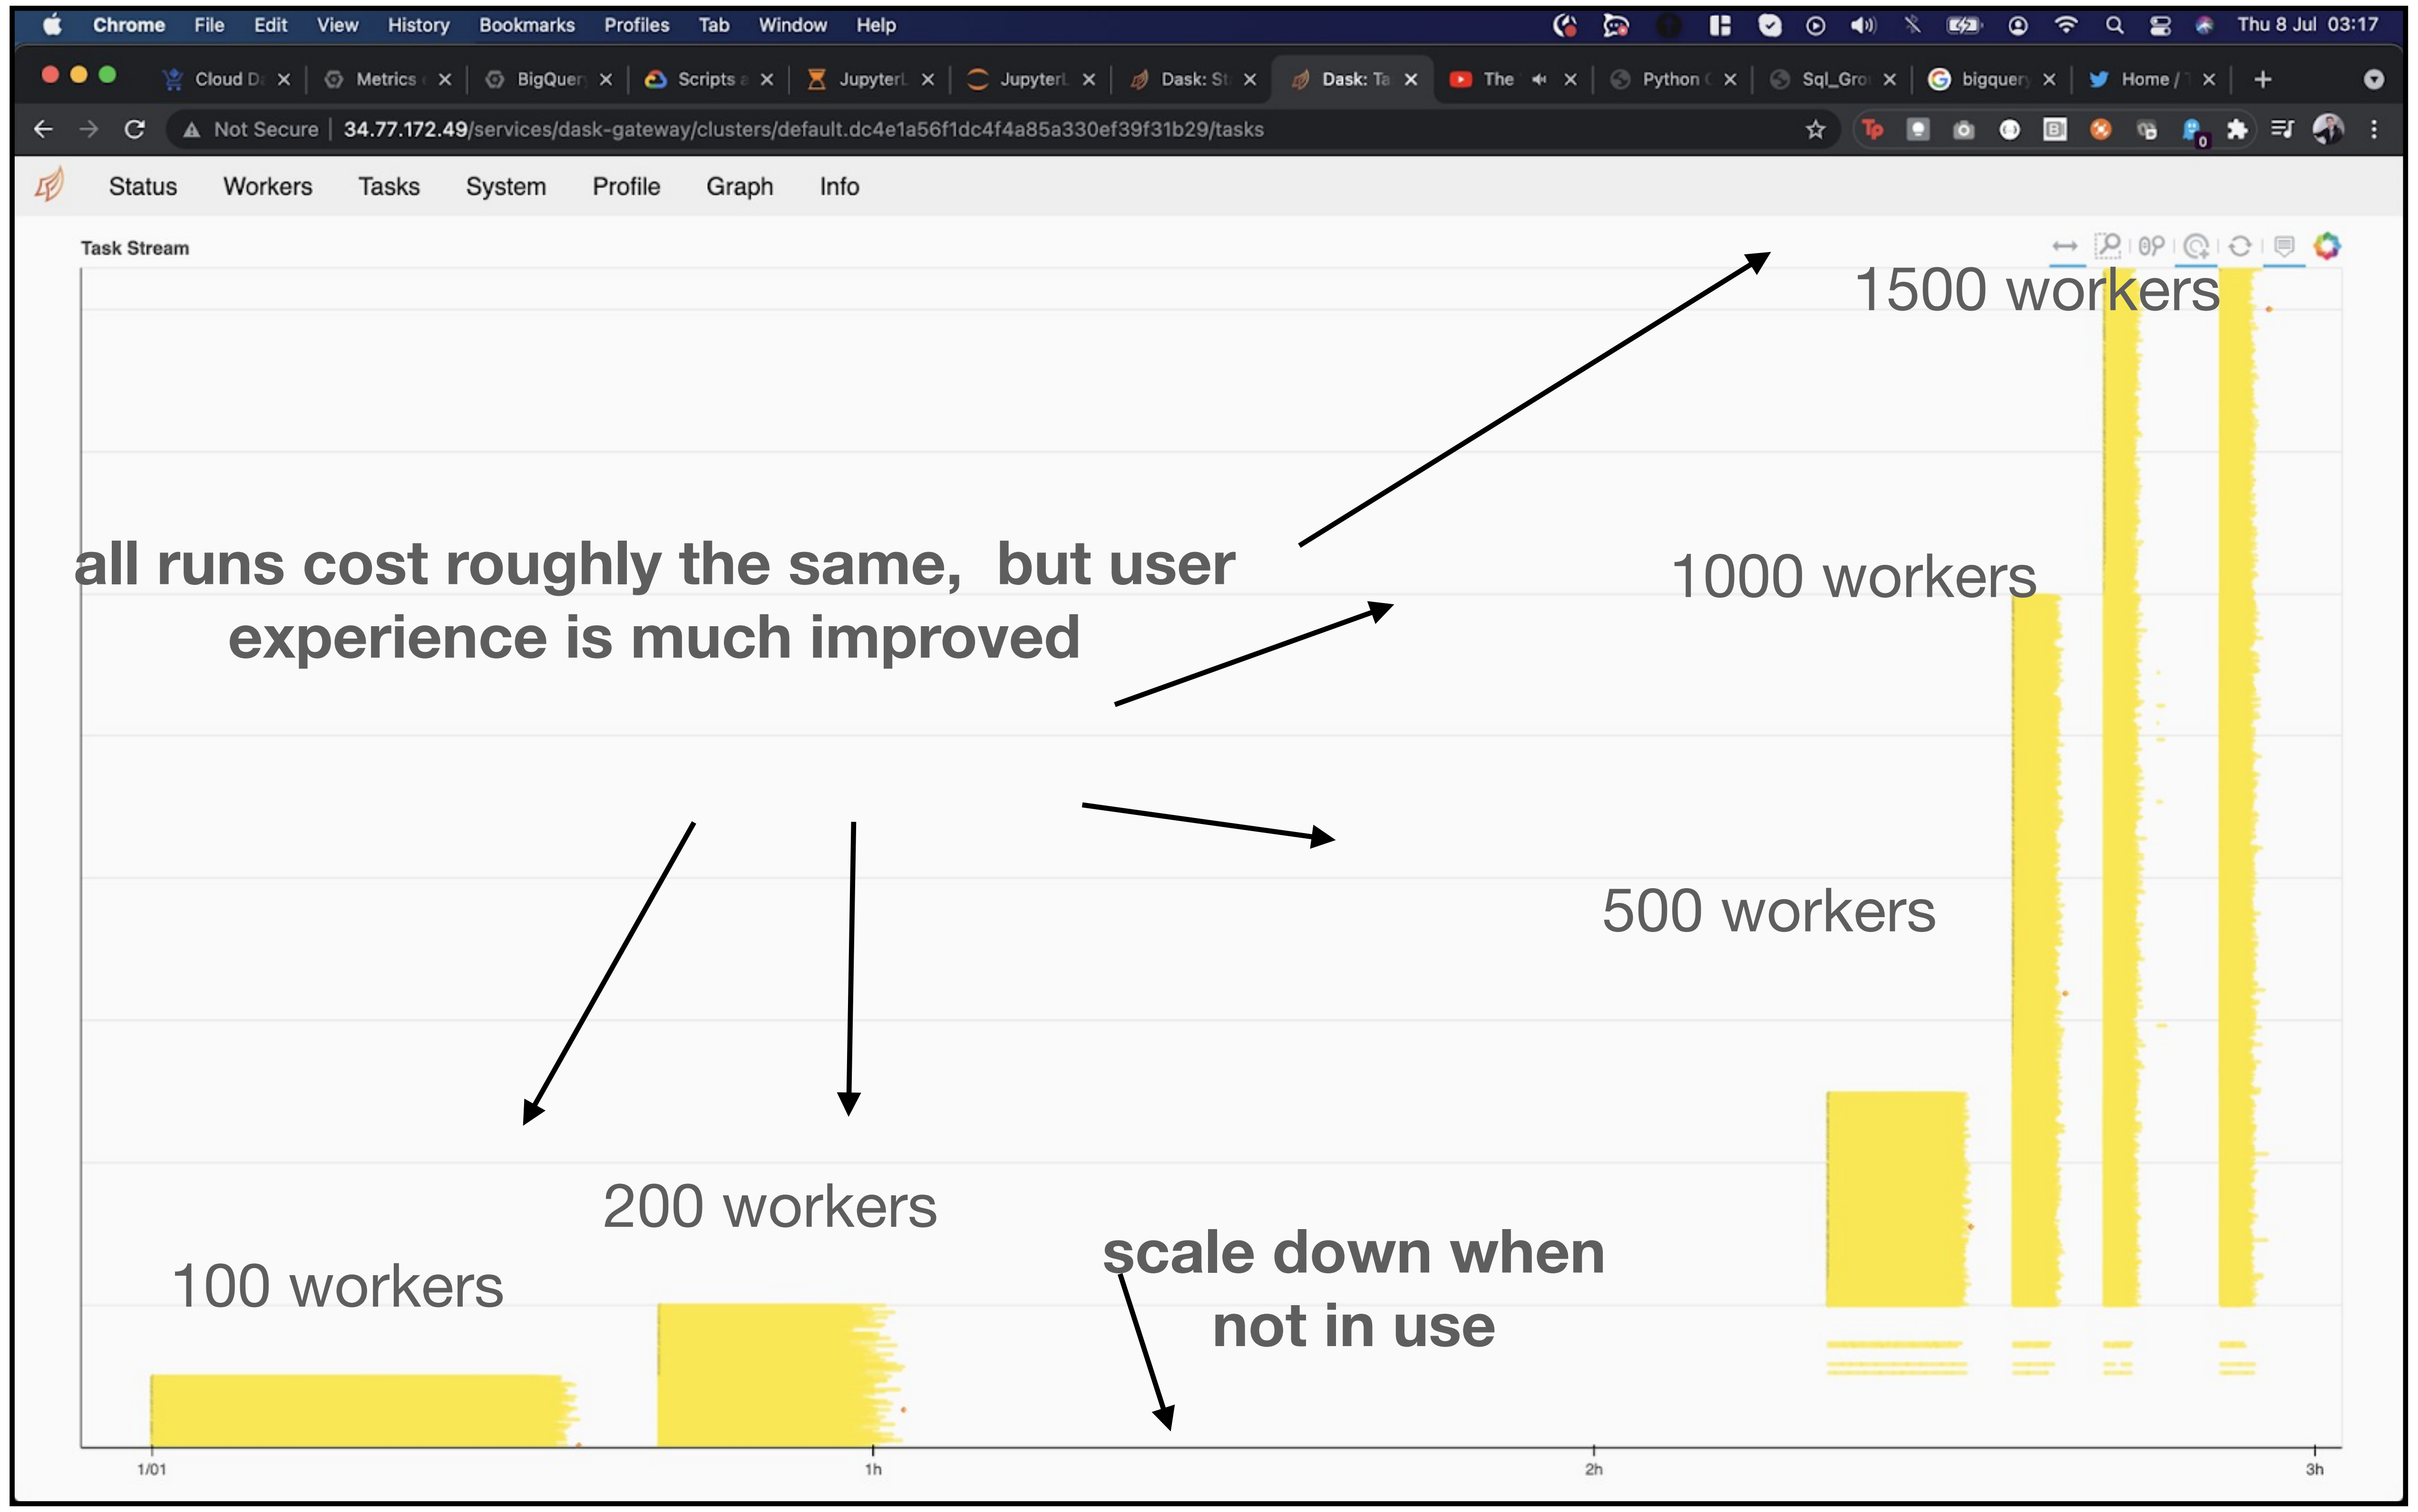

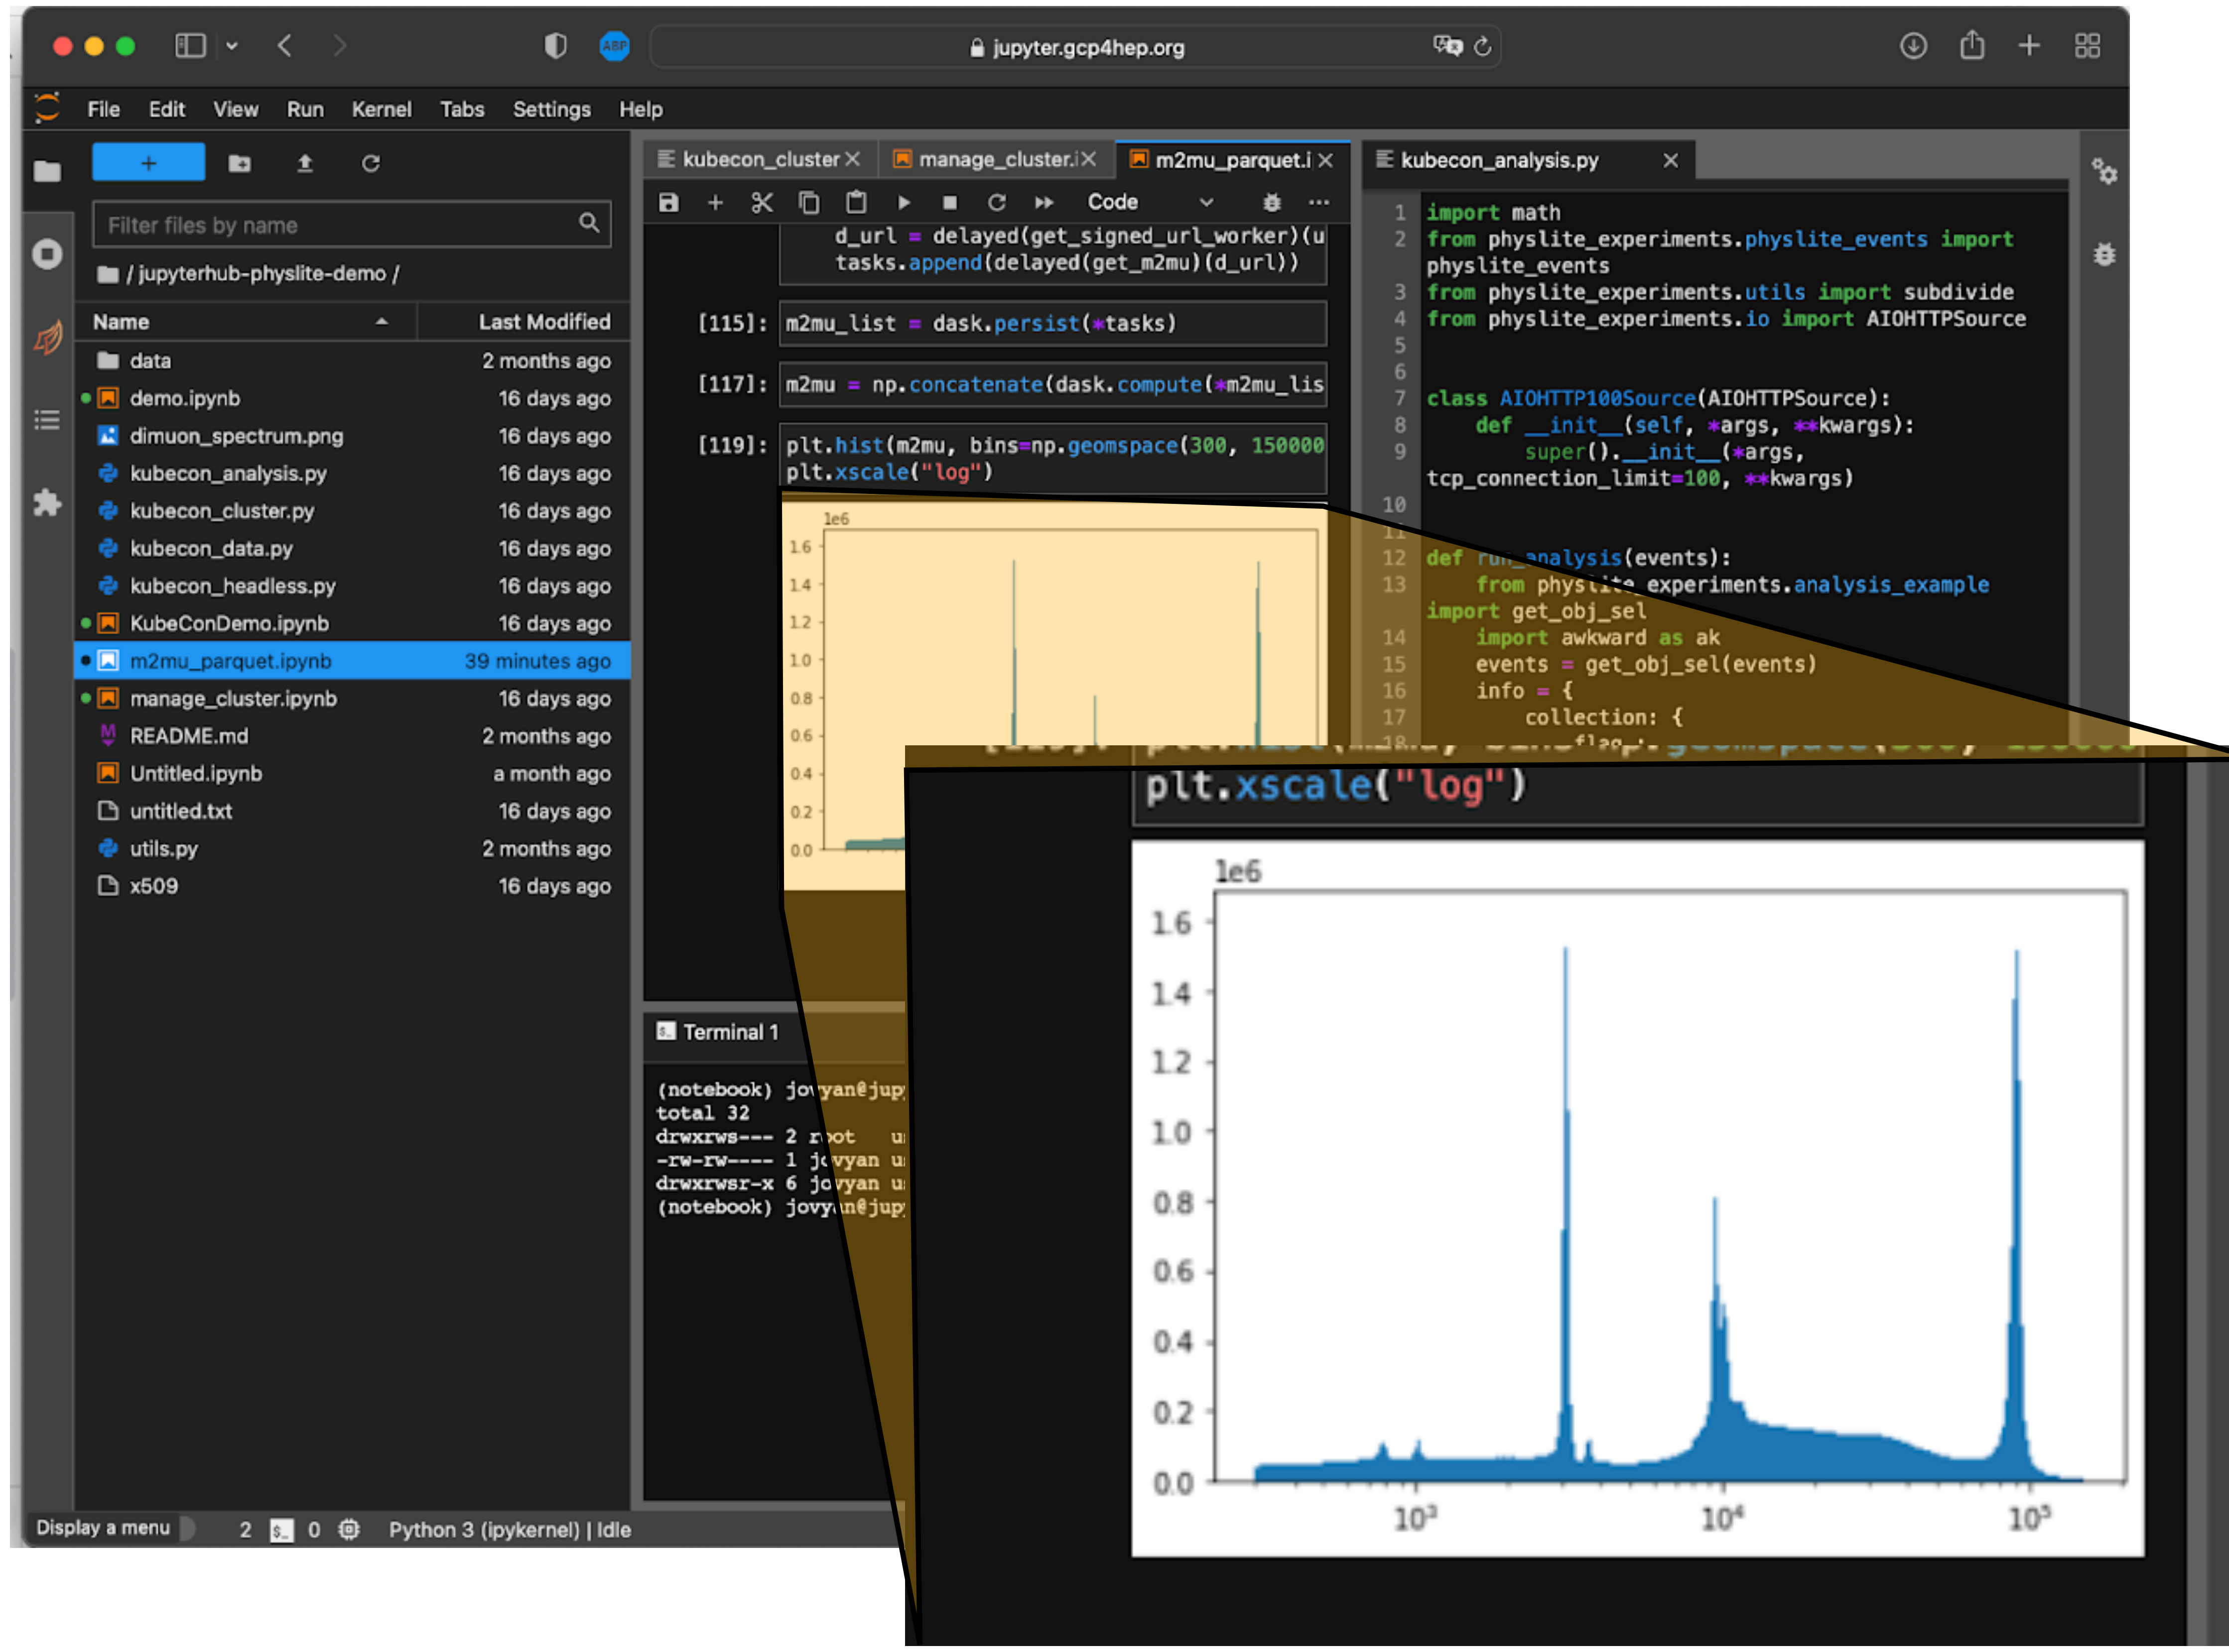

Wrapping Up

We have all the pieces together to build future analysis facilities & know we can scale: Opportunity for close collaboration across communities

“Forschungsdateninfrastruktur!”

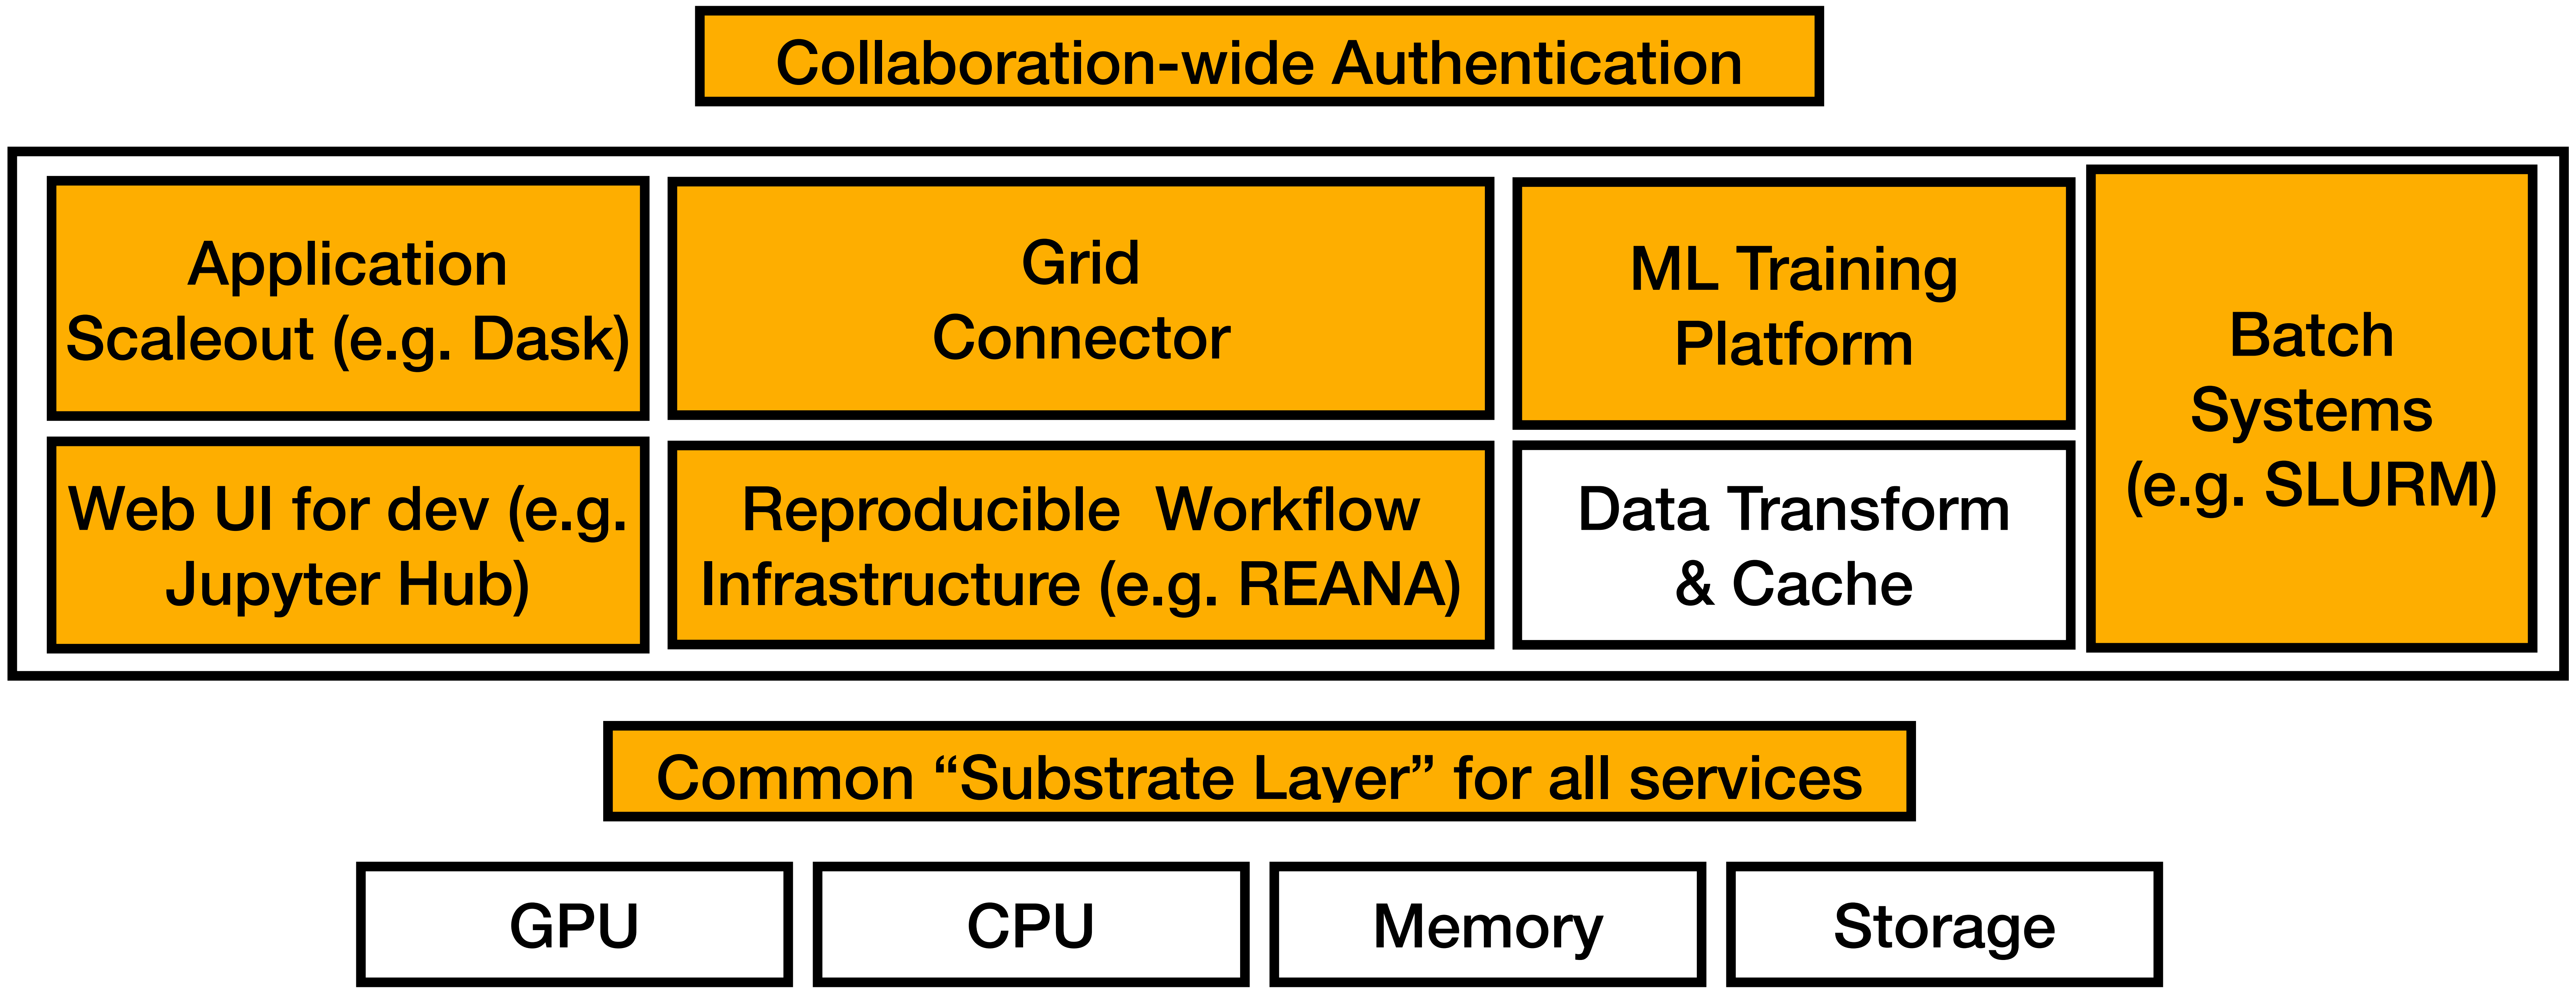

Wrapping Up

Kubernetes seems to be a great option, but we need to work more on our ability to deploy it within an academic context (HPC, Labs, Universities) → question of policy and training

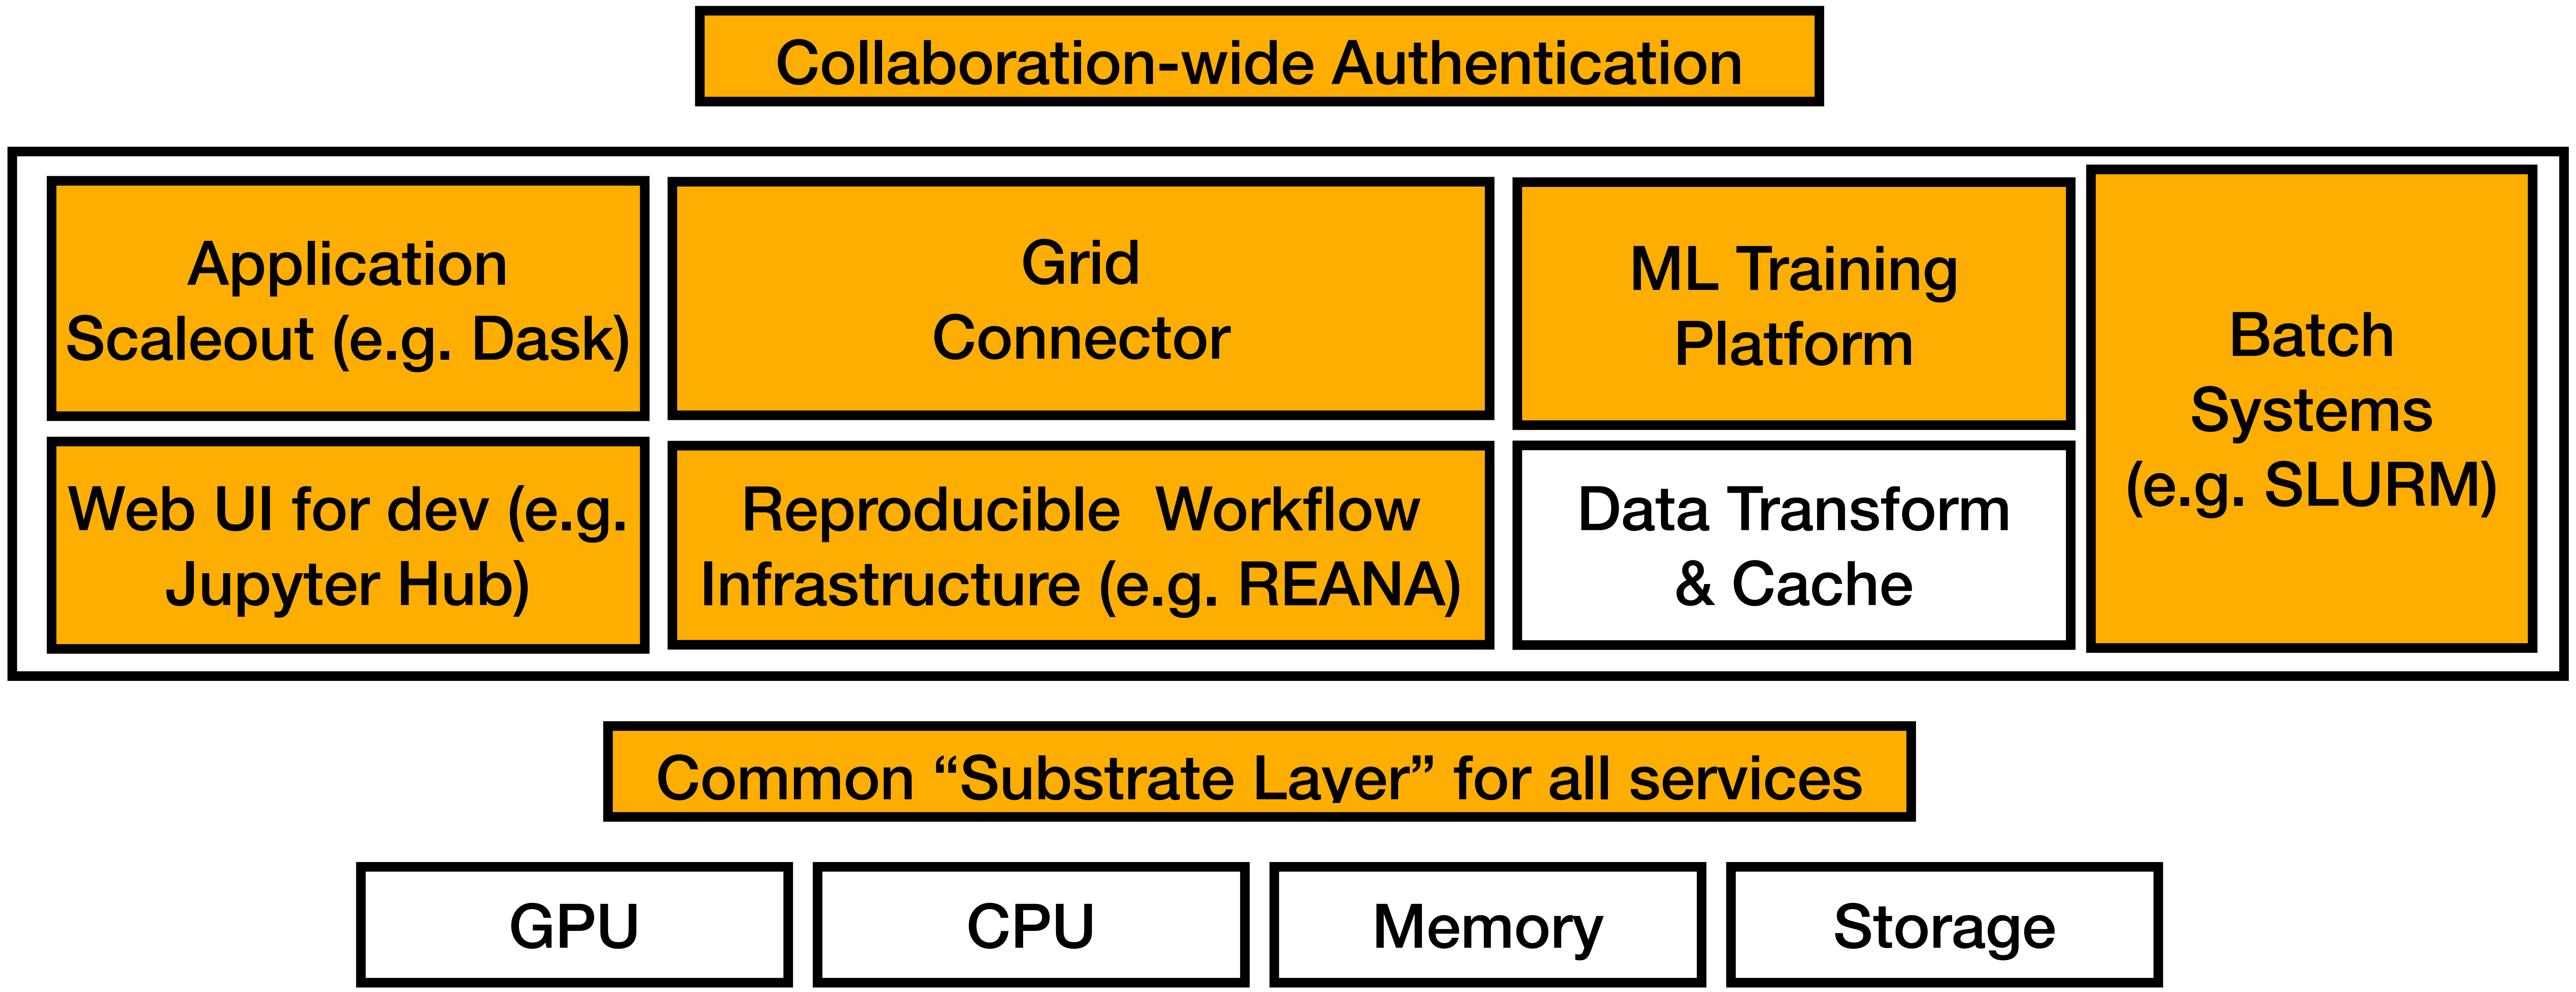

Wrapping Up

HEP Software Foundation is providing a Forum to exchange ideas, experience as we work towards this future

Thursdays at 17:00
biweekly meetings