The Case for a “Big Church” AF

Paolo Calafiura (LBNL)

The Case for a “Big Church” AF: scalability

US ATLAS survey Oct 2020
Adelman, Hance, Martinez Outschoorn

- Why 75% of users produces their derived data (n-tuples) on the grid?

My guess

- *To be able to run at scale using centrally managed resources and tools*

The Case for a “Big Church” AF: community

- Why are 50% of US ATLAS physicist still running on overcrowded lxplus?
  - *To share readily **data**, **know-how***

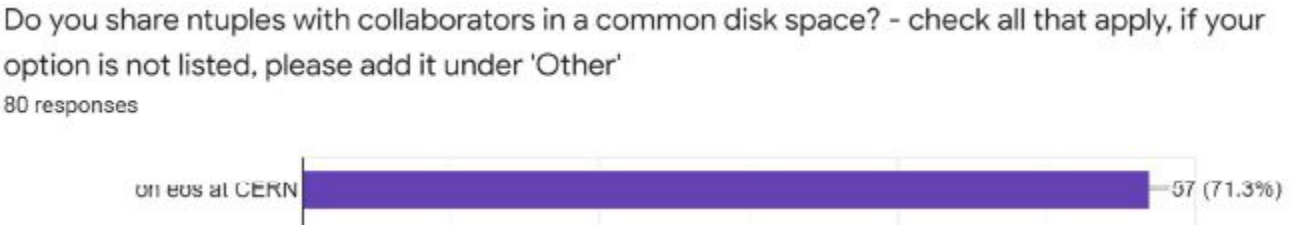

The Case for a “Big Church” AF: usability

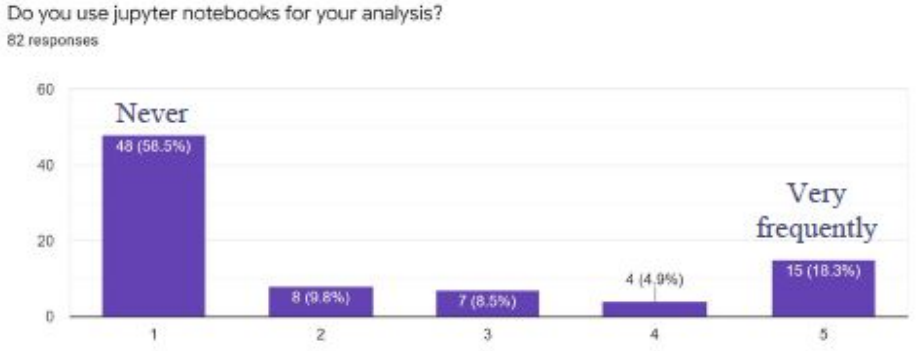

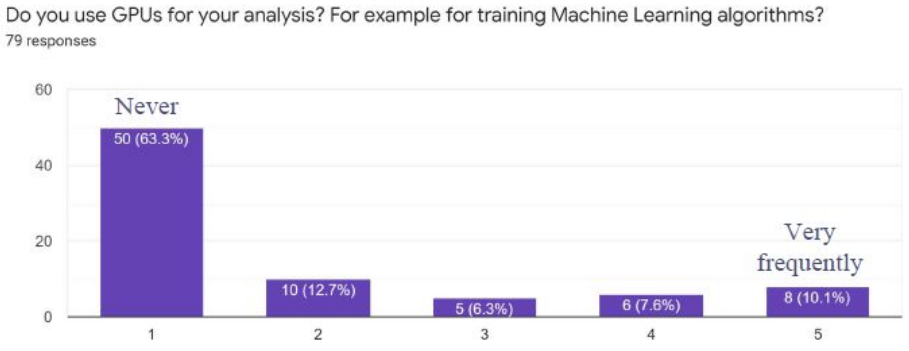

- Why 60% of users never used a jupyter notebook and (presumably) the scientific python stack?
- Why <20% run regularly on powerful, underutilized GPU resources?

*Users have no time to learn how to be efficient: they **stick to defaults***
- *Migrating an analysis group is even harder: **shared knowledge***

The limits of Shared T3s

Each US ATLAS Shared T3 has ~50 users and ~0.5 FTE dedicated to user support

For comparison, NERSC has ~7000 users and ~15 FTEs in two user engagement groups

- While ratio is favorable to US ATLAS, the problem is the lack of **critical mass**
  - Can't share support load,
  - Hard to have specialized support e.g. in data formats, vs ML tools
  - No **career path** for user support personnel within an experiment shared T3

Each US ATLAS Shared T3 offers plenty of resources to its users
but, ultimately, it is an "island"

- Not straightforward to share data, reproduce results outside the shared T3
- Even the largest shared T3 is tiny compared to ATLAS grid resources.

Support is the Key

If we want AFs to succeed, we need to overwhelm users with support:

- User-level support:
  - From beginner defaults, to cluster configurability
    - We are doing this, need more effort, particularly in entry-level tutorials
- Embedded support:
  - Analysis groups are high-value customers, embed one “**account executive**”
    - Ideally a **50/50 physics/tools** person with knowledge & contacts on both sides
      - These people are precious but do exist, we need a **career path** for them
- DevOps support:
  - Cutting edge analysis tools are hard to deploy at scale
    - Fledgling **DevOps community**
      - Highly marketable skills, need to create a **pipeline** and a rewarding career path

A Big Church Analysis Facility

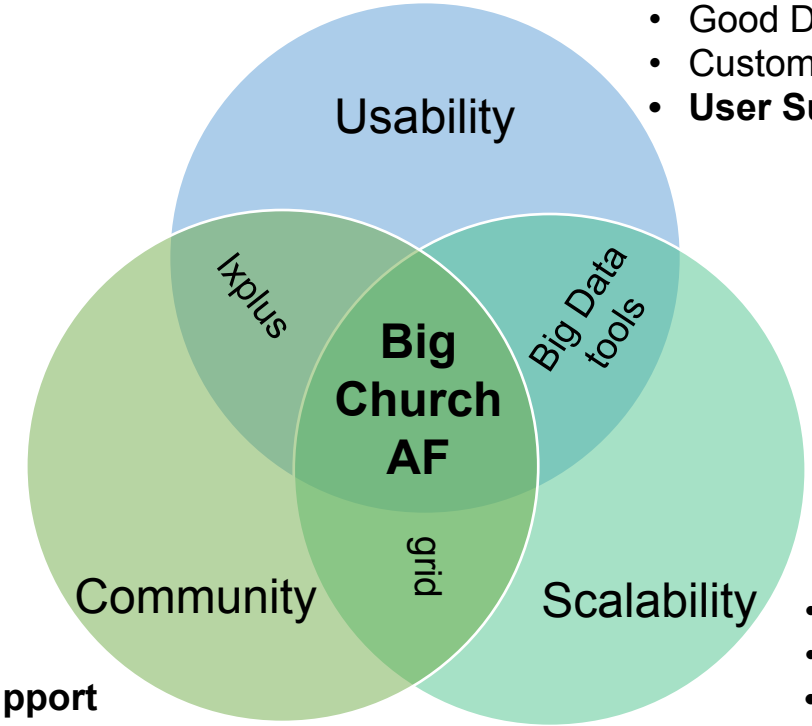

- Good Defaults
- Customization
- **User Support**

- Reusability
- Reproducibility
- **Embedded Support**

- Data Parallelism
- Task Parallelism
- **DevOps Support**

Who would pay for this?

The HPC Model

- Funding agencies pay upfront for all costs, and allocate support and hardware resources

The Cloud Model

- Funding agencies (DOE, NSF or both) provide
  - hardware infrastructure (power, cooling, networking, etc)
  - management, sysadmin and core devops effort,
  - training/education programs, including fellowships etc
- Science groups buy into the AF by
  - Paying for hardware resources they use either on a pay-as-you-go basis or long-term (e.g. for storage).
  - Funding (fractions of) devops and user support personnel, and providing embedded support to their analysis groups.

With any model, support structure will be key

The Hybrid Model

- A combination of the two above
  - Groups with long-term commitment buy into the program
  - Smaller, start-up groups rely on an allocation of resources and centralized support

Possible next steps

Define scope and requirements

- HL-LHC experiments? Cross-frontiers? All data intensive sciences?
- End-user analysis?
- All data-intensive workflows?

Estimate budget, identify sponsors, determine funding model

Define strawman architecture(s)

- Single-site? Distributed? Cloud?
- Common computing platform? Menu of user-selectable configurations? BYO hardware?
- Storage/data management model (hard to share storage resources)

Snowmass CompF4 and CompF5 working groups could host these discussions (Snowmass LOI)