The Evolution of HEP Analysis Facilities

Paolo Calafiura <pcalafiura@lbl.gov>
V0, 8/31/2020

Abstract

This LOI will review the role of Analysis Facilities (AF) in HEP collaborations. It will then present the requirements AFs must deliver on cost-effectiveness, community building, scalability, and usability, and argue that the key to AFs success is to focus on the support model for the analysis communities.

Introduction

Many HEP collaborations rely on a distributed, tiered computing model for their data processing and storage needs. The bulk of HEP grid resources remain dedicated to high throughput computing and large, centralized tasks, including the production of Analysis Object Data (AOD), and multiple streams of derived AODs (dAOD). While the tiered computing model developed twenty years ago by the MONARC group has become considerably less hierarchical over time, analysis workflows are still run by individual physicists or physics groups on dedicated resources, carefully insulated from production workflows. In what follows, analysis workflows are defined as[1]

1. derivation of columnar data ("n-tuples") from centrally produced [d]AOD,
2. interactive analysis of these datasets to produce physics results (plots, fits, etc).
3. Fast, particle-level simulations to study the impact of systematic effects on physics results

Analysis Facilities in 2020

Analysis workflows are often run at institutional Tier 3s, locally owned clusters operated by one university or lab. The main strengths and weaknesses of an institutional T3 derive from the fact that it serves a small user community. Users have direct control of their resources, which are often underutilized. User support and system maintenance are typically provided by fellow group members, who are easily accessible and deeply understand user needs but may not have the time or know-how to maintain the T3 professionally.

To address these issues, large HEP collaborations have introduced Analysis Facilities (also known as Shared T3s, or Regional Analysis Centers). AFs are often co-located, and share resources with T1 or T2 grid facilities, but are designed to support analysis workflows and the physicists running them. They provide interactive logins, local batch capabilities, user home directories, and group or individual data storage. They are professionally maintained. Being integrated into their host grid site, they provide reliable access to AOD datasets and quick turnaround times. They support both the traditional ROOT analysis platform, and the ML/scientific python/jupyter platform.

[1] While this is not a complete list, it attempts to capture the main characteristics of analysis workflows.

Unfortunately, at least in the case of US ATLAS, AFs are not as widely utilized as one may expect: for example, a recent survey of US ATLAS physicists indicated that 64% do not utilize the two US ATLAS AFs, although 60% of the respondents indicated they expect to run their analysis workflows on the AFs five years from now. On the other hand, most US ATLAS physicists rely on the heavily oversubscribed CERN lxplus cluster for at least some aspects of their analysis. While counterintuitive, this is likely a manifestation of the fact that physicists operate in a highly competitive environment and have no time to explore new and more efficient ways to run their analysis workflows. They are provided defaults and stick to those. Even when users are willing to explore new paradigms, they still have to convince a whole community (e.g. an analysis group) to migrate, before they can fully take advantage of them.

The Key Role of Support

This LOI argues that support is key to the success of an Analysis Facility. In fact, if we want AFs to succeed, we need to overwhelm users with right kind the support:

User-level support:

from providing beginners with accounts, default settings, and tutorials, to enabling advanced users to define their analysis platform.

Existing US facilities already provide administrative support to their users, and the larger collaborations provide experiment-specific tool support through their Operations programs. but more effort is needed as the AF user base grows.

Embedded support:

Physics Analysis groups ranging from a handful to hundreds of users are high-value customers. AFs should embed one "account executive" per group, ideally a 50/50 person with knowledge and contacts on both sides.

These crossover experts are precious but do exist. AFs urgently need to provide a career path for them.

DevOps support:

Cutting edge analysis tools are hard to deploy at scale. DevOps are needed to tailor existing tools to the AF platform specifics and the community needs.

There is a fledgling DevOps community in HEP. DevOps have highly marketable skills; hence the AFs need to create a pipeline to train them and a rewarding career path for those who chose to remain.

Analysis Facilities in 2030

HEP experiments are faced with an order of magnitude increase in data volumes, and similar increases in data complexity. To maximize efficiency, resources will have to be pooled, and workflows scaled up opportunistically. The additional complexity of analysis tools will affect end-users, particularly those at institutional T3s. Even larger collaborations running shared AFs will have difficulties providing the extra support to their physics analysis groups.

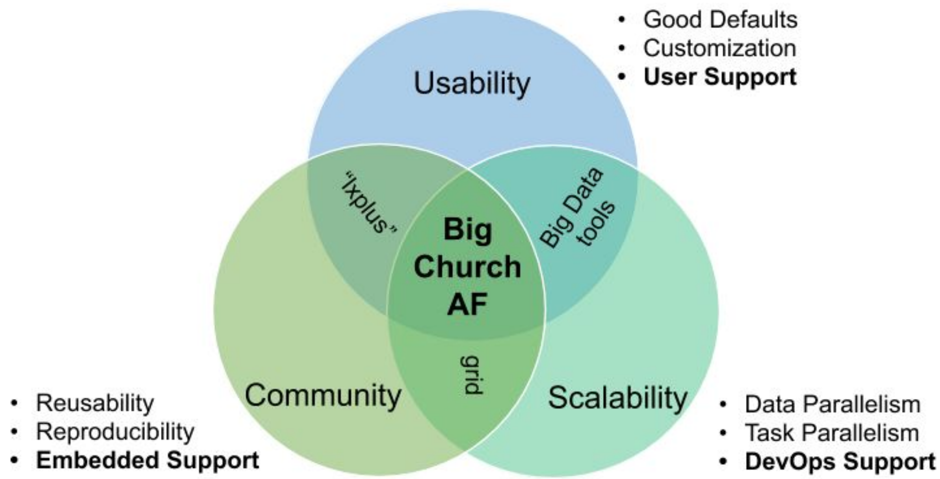

Fig 1: A Big Church Analysis Facility

US HPC centers are very good at building and supporting HPC science communities, and at training and creating career paths for experts (often former domain scientists) willing to support them.

The question this LOI proposes to explore is how can the HEP community replicate the successful HPC science support model while addressing the unique requirements of HEP physics analysis[2]. Besides the fiscal challenges, individual experiments may find it difficult to attract and retain experts, e.g., in distributed ML frameworks. Management and funding agencies may reasonably argue that the expertise needed to support physics analysis is largely experiment-independent, and should be pooled, just like software R&D or hardware resources. At a very high level, Fig.1 sketches a Big Church Analysis Facility which would serve the entire US HEP community, and potentially other "big data" science communities.

[2] Unfortunately, the core mission of HPC centers is to provide resources to solve computational problems at extreme scales, not to support thousands of physicists analyzing PBs of n-tuples.