ATLAS Software and Computing: R&D program in 5mins

CERN IT R&D advisory group meeting

22 September 2021

Zach Marshall (LBNL) and Alessandro Di Girolamo (CERN)

The ATLAS Computing Conceptual Design Report

- The ATLAS CDR, LHCC-G-178 public document, was submitted to the LHCC in May 2020
  - Contains the ATLAS Software and Computing strategy towards HL-LHC
- ATLAS S&C community now working on a second document, complementing the CDR, which will describe more in detail the milestones, deliverables and timeline of the various projects
  - To be submitted to the LHCC beginning of October
- Summarize in 5 minutes the rich R&D projects is 

  - But we can maybe discuss ways to engage the community on the challenges

R&D projects

- One of the most critical parts of our environment is our community
  - “Anything is possible”, if there are the right persons with the right attitude
  - We have a proven record of making things happen thanks to superheroes
  - Fantastic skills and great motivations are definitely a huge asset in our sleeves
- We need to make sure that our community is sustainable
  - Right now we are really not sure about our sustainability

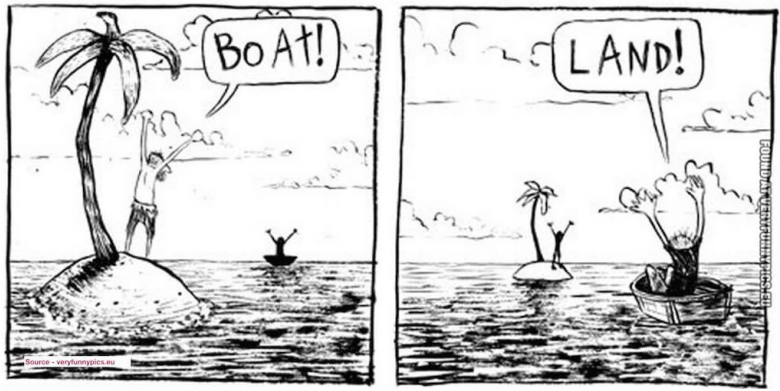

- We do not have a problem today, but we might have one in the (not so far) future
  - Possible Q: Why we don’t have a problem today?
  - Possible A: Thanks to the great work several years ago to build this environment, and our community.
- R&D projects are a fantastic way to engage more people
  - and also to keep our experts excited and challenged
  - We need to make sure we are able to integrate the “new” people into our community, as we need to make sure that we are able to integrate the successful R&D into our experiments

Presented LHCC meeting 31 August 2021

ATLAS - Personpower concerns

Presented LHCC meeting 1 June 2021- Still accurate

Personpo

- Stretched to the limit in computing operations, concerned for a successful Run 3
  - Can essentially only get by right now, but expect a more intense period starting in 2022
  - Desperately struggling to find experts in several areas.
- We (ATLAS S&C) cannot able to offer "real" jobs: only temporary support
  - Recent changes in CERN users conditions makes it even less attractive, and risk to demotivate people
    - We rely on people's passion and motivation!
  - We need concrete career paths planned for Software and Computing professional from the FAs and the Institutes.
- Some FAs invest into (needed) R&D but without the commitment to maintenance and operations afterwards
  - Strategically, we want to have groups evenly participating in R&D projects and operations. Beneficial for the effectiveness of the project and for the sustainability.
- Some FAs are also having troubles in maintaining (some of) the commitments taken because of expertise leaving
- We have started F2F chats with each ATLAS ICB (International Computing Board, representative of the Funding Agencies) member:
  - Very good return in terms of people understanding and sharing our concerns
  - But the overall lack of concrete possibilities remains concerning

Ale Di Girolamo & Zach Marshall - 1 June 2021

16

ATLAS - Personpower concerns - more details

Stretched to the limit in computing operations, concerned for a successful Run 3

- Notable examples are:
  - Distributed Computing operations:
    - Dealing with the daily failures of our very efficient distributed infrastructure.
    - Since beg of 2021, ~350 tickets to sites, plus order of 4k emails about issues and/or optimizations (both ATLAS – e.g. Rucio/Panda/Physicists expectations – and CERN services related, e.g. monitoring, FTS)
  - Infrastructure, e.g.
    - Build systems
      - CI and nightlies requires almost daily attention
    - Central Services
      - ATLAS Central Service team manages >600 nodes(VM) with >150 different services
  - Integration of new versions of tools/integration of new tools
    - New G4 versions, new Generators (e.g. MadGraph, Pythia) versions, new ROOT versions

ATLAS - Personpower concerns - more together?

- Distributed Computing operations:
  - Sharing of, and dealing with, issues with other experiments: we use (mostly) the same infrastructure (network, shared sites) and many similar frameworks and tools (FTS, Rucio, CVMFS).
  - More and better automation – deploying ideas, not just R&D towards automation
- Integration:
  - Boundaries and responsibilities are (reasonably) well defined between experiments and both the various projects like Root, G4, various generators and infrastructure work like e.g. building SW on heterogeneous resources (e.g. other architectures, GPU)
  - Still, having experts working across the boundaries, deeply involved in both the projects and the experiments could be tremendously beneficial for all of us
    - Technical validation of the new versions in real ATLAS conditions could be very useful

A support model with persons participating in experiments DevOps (and integration) for a concrete part of their time could be advantageous:

- Just part of the job: bring the project/service/framework dev view to the experiments and take back to the project the experiments requirements and experiences

A few ideas

- QoS, DataPopularity
  - understand better (and thus optimize) the usage of our storage
- ARM, Power9
  - HW and support for (e.g. LCG) builds and CI; validation of common software (Generators, Geant4, etc)
- DB:
  - (more on the D than the R of the R&D): e.g. CREST
- Artificial Intelligence, ML
  - Lot of activity already ongoing (e.g. ATLAS ML workshop April 2021, 50+ talks, Operational Intelligence)
  - (one of the) challenge is getting concrete impact out
- Software (re)engineering
  - A Common Tracking Software (ACTS), Rucio/FTS
  - Heterogeneous resources, specifically GPU:
    - workload offloading, I/O, memory management (n.b. ATLAS uses Gaudi, like LHCb)
    - e.g.: madgraph generator (discussed already a few times - negative weight is also a concrete challenge)

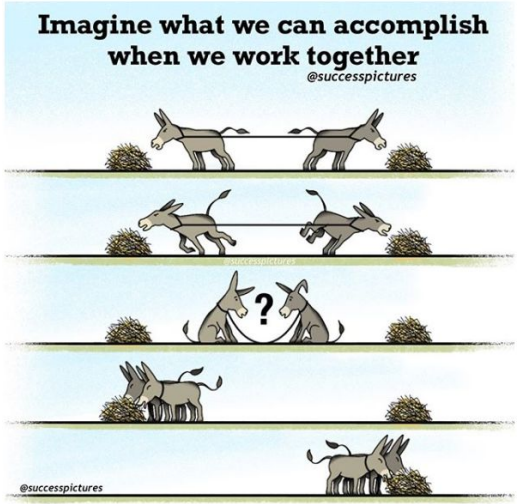

Next

- We are evolving the ATLAS strategy described in the CDR
  - Impact and effort needed (how much, skills) being detailed
- We are stretched thin in several areas
  - SuperExperts are just a few.
  - Evolving the experiments SW and C to make it more modular is a huge challenge,[very difficult to evaluate the possible positive impact]
- Persons' skills are not all equal, we need an agile and iterative way to engage in the HL-LHC challenges depending on the possibilities
  - We would obviously benefit from more skilled people, but we know there is no free lunch:
  - We might need to decide on which areas/projects we can have critical mass instead of just sharing evenly ourselves on everything.
- Would be happy to keep on engaging with CERN (IT, EP, and more) to iteratively define the possibilities and priorities