CHEP 2018 Highlights

Michael Davis

Sofia, Bulgaria

Sofia, Bulgaria

Sofia, Bulgaria

Palace of Culture

Sofia, Bulgaria

Palace of Culture

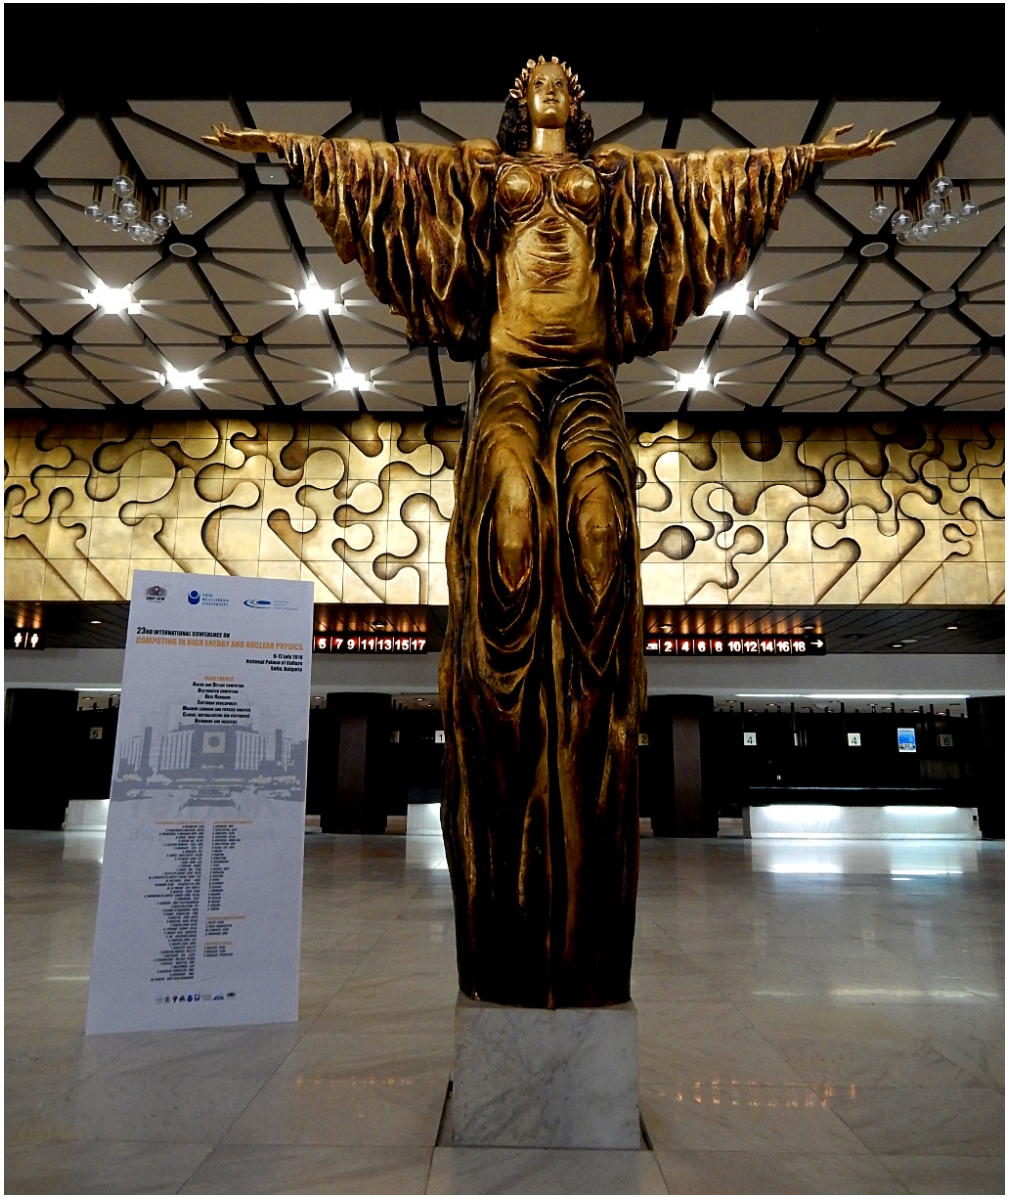

Sofia, Bulgaria

Palace of Culture

Looking forward to Run-3 and beyond

Motivation

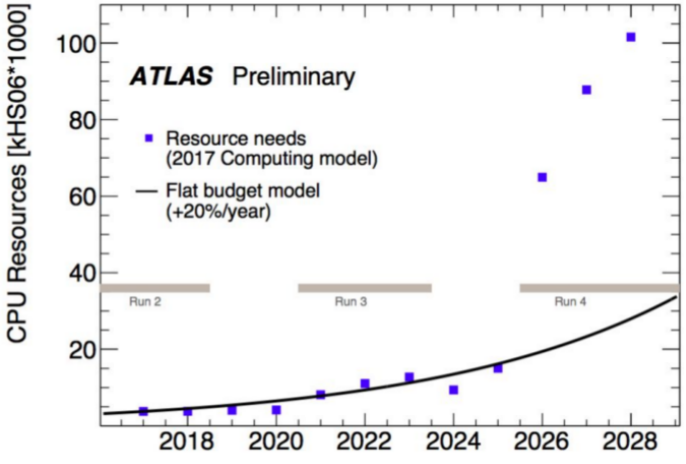

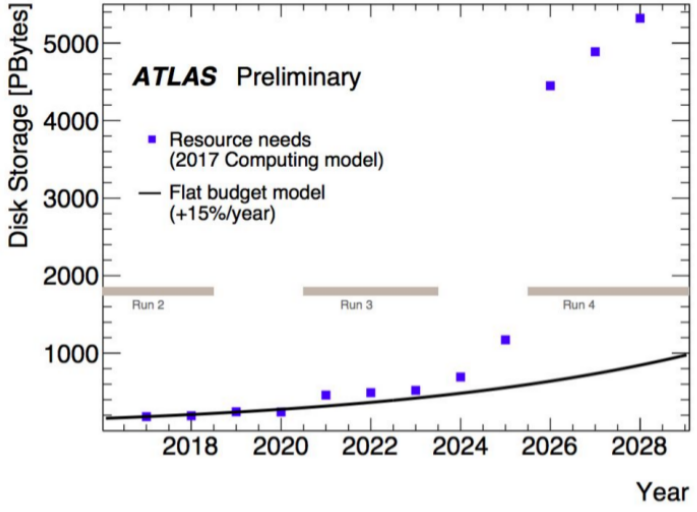

- LHC computing needs keep increasing, while budget is flat at best
- IT landscapes, computing infrastructures and funding models change
- Heterogeneous workloads, architectures, resource types, storages
- We need to be able to use every resource available and use it efficiently
- ...and there is a general manpower limitation

Looking forward to Run-3 and beyond

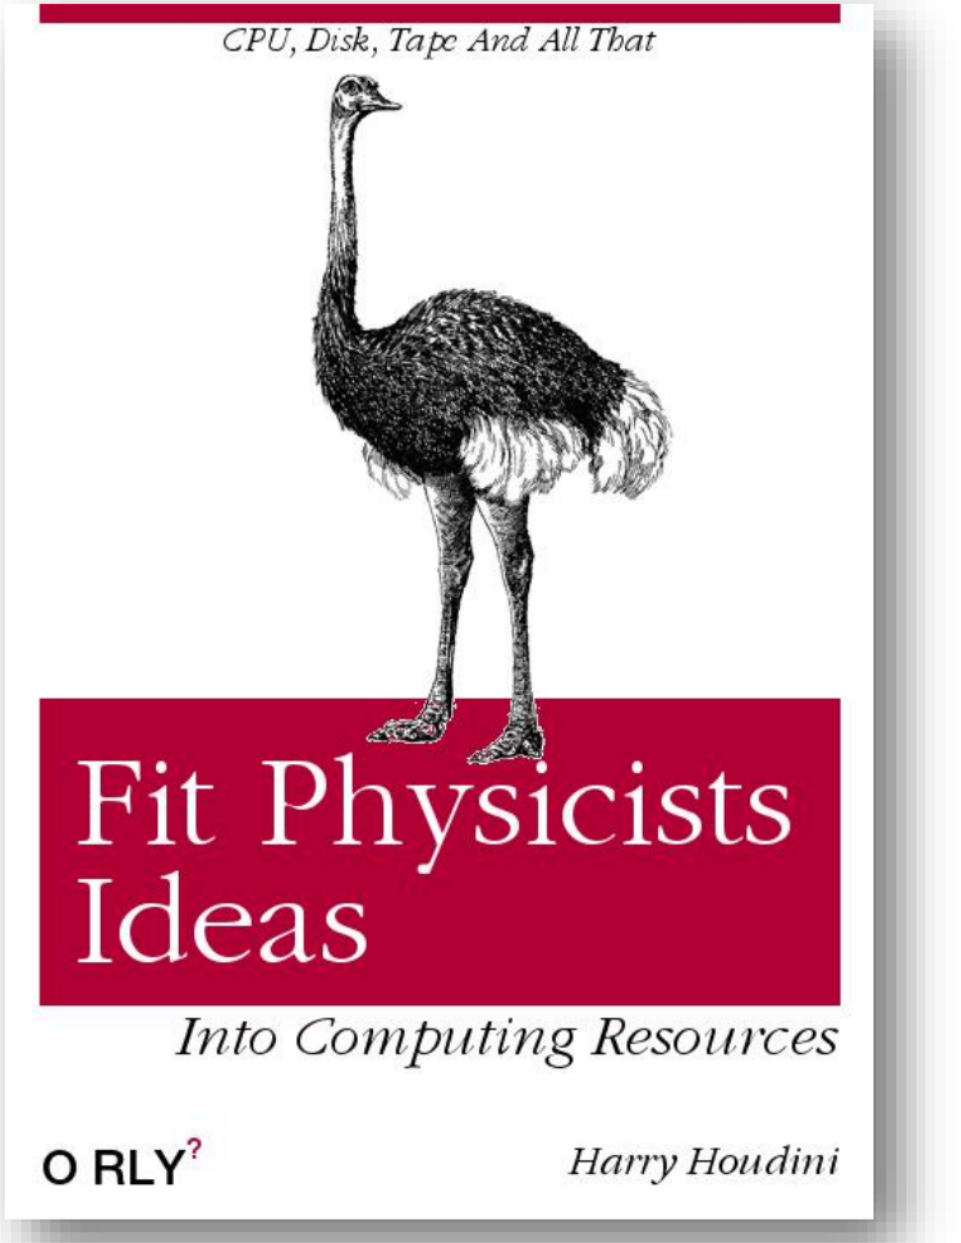

Looking forward to Run-3 and beyond

LHCb Trigger

Today

Run 3

Looking forward to Run-3 and beyond

Storage Requirements

- Storage needs are driven by data of HLT output bandwidth
  - Tape needs incompressible, while mitigations possible for disk
    - E.g. parking scenarios are considered but introduce additional operational costs for the experiment and infrastructure costs for sites
- MC simulation output data format mostly migrated to m(icro)DST format with small contribution to needs introducing a size reduction of factor 20

- LHCb relies on a small amount of sites with disk storage:
  - T0 + 7 T1s + 13 T2s with minimum size requirements especially for T2s
  - Data caching especially on ”small disk sites” is not a major use case

HNSciCloud

Helix Nebula Hybrid Cloud Model

Bringing together:

- Research Organisations
- Data Providers
- Publicly funded e-infrastructures
- Commercial cloud providers

with:

Procurement and Governance suitable for the dynamic cloud market

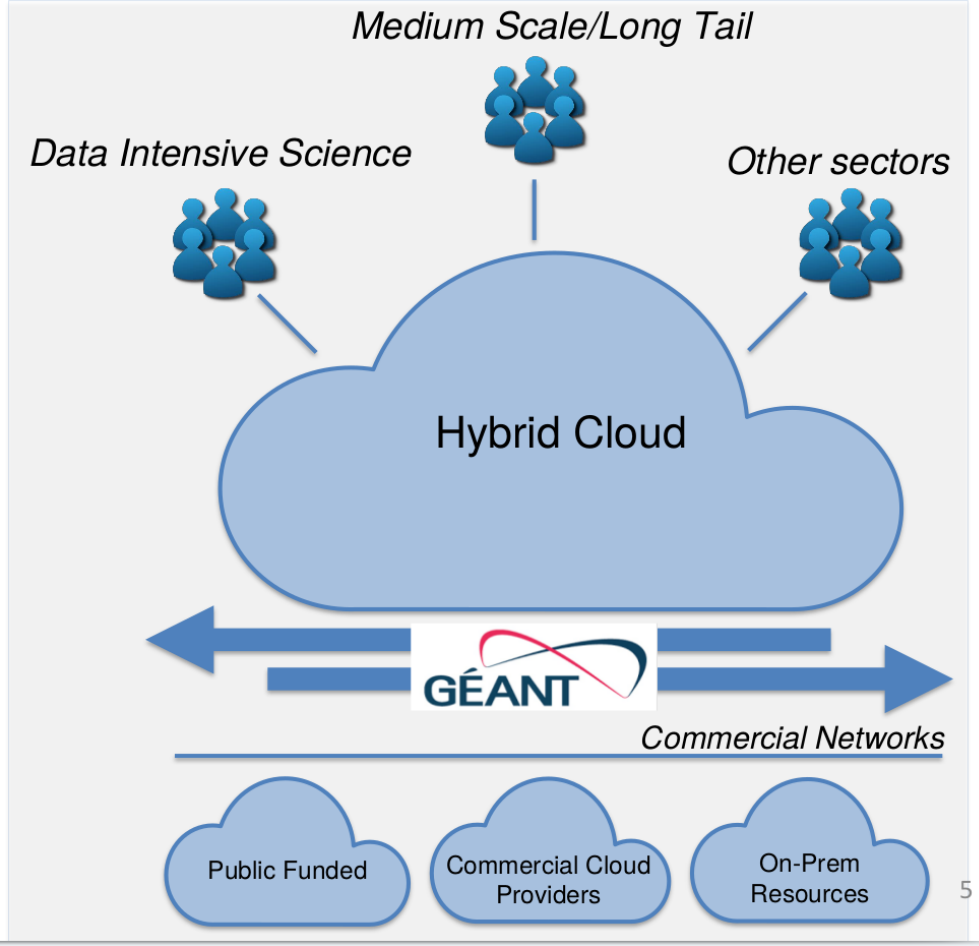

HNSciCloud

HNSciCloud project phases

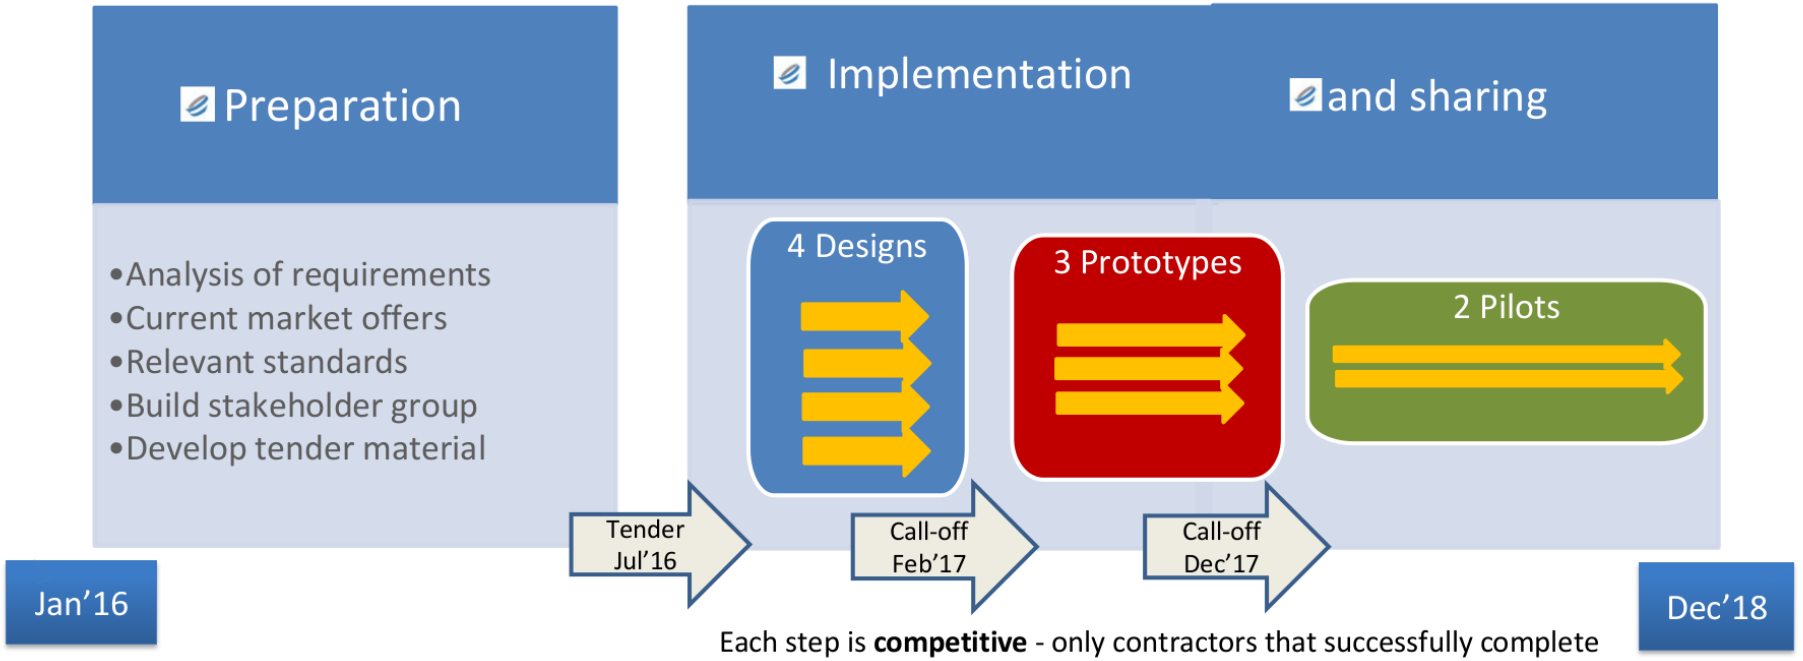

Each step is **competitive** - only contractors that successfully complete the previous step can bid in the next

Phases of the tender are defined by the Horizon 2020 Pre-Commercial Procurement financial instrument

HNSciCloud

Cloud Providers

- *T-Systems*
  - *IaaS based on OTC*

- *RHEA*
  - *IaaS provided by Exoscale*

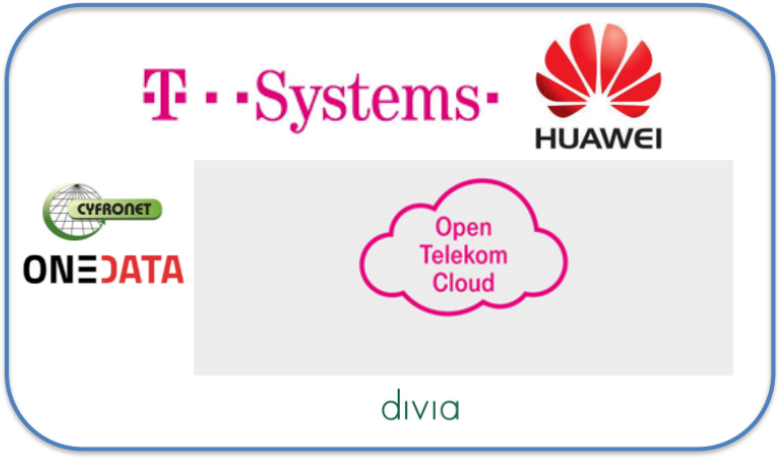

Data Management Frameworks

ATLAS DISTRIBUTED COMPUTING OVERVIEW

The ATLAS distributed computing system is centered around:

- **Workflow management system**: Panda (Talk 143)
- **Data management system**: Rucio (Talk 137)
- Many **additional components**: AGIS, ProdSys, Analytics, ...
- **Resources**: WLCG grid sites, Tier0, HPCs, Boinc, Cloud

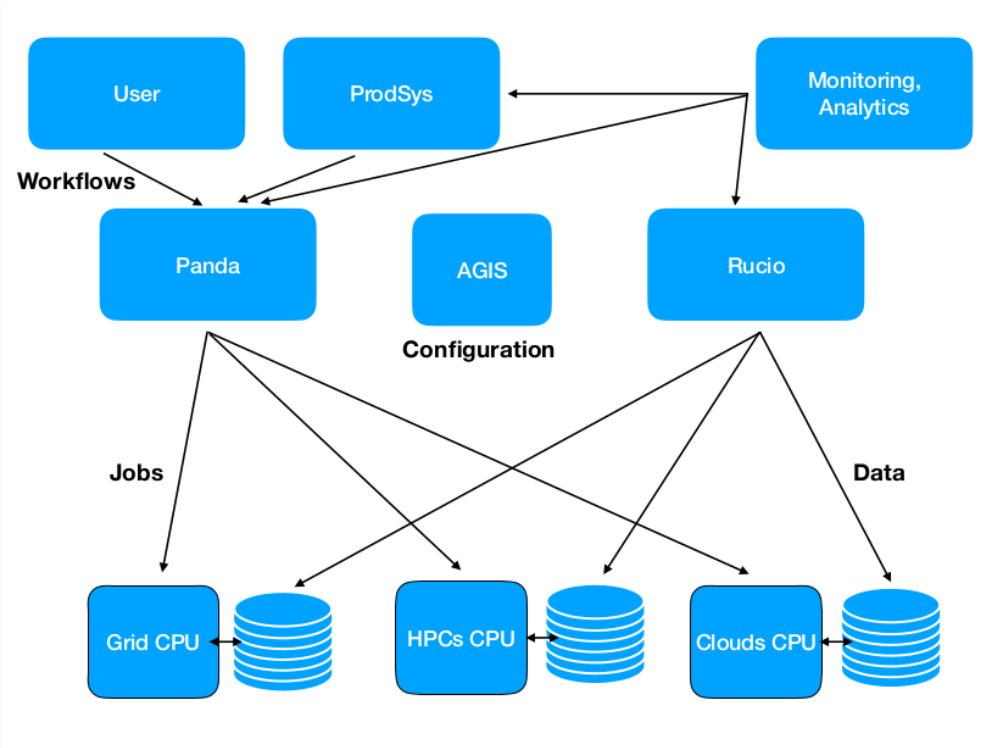

More details: Poster 144, Poster 141

Rucio (ATLAS)

Evolution of the Systems - Data management & Site infrastructure

Rucio: (Talk 137)

- Interest in Rucio by other HEP experiments and communities - Very successful Rucio Community Workshop in March
- Examples of new features: Adding rucio input file mover to panda pilot, site file cache awareness (Poster 138), object stores (Poster 162), dynamic data placement (Talk 140) zip archive file creation, Tape carousel

Data lake/ocean:

- A very successful R&D project with Google: integration with Rucio and distributed analysis using Harvester (Talk 133)
- Ramping up in the WLCG DOMA (data organisation, management, access) project to explore possibilities to overcome disk shortage in HL-LHC

Singularity and Containers: (Poster 163)

Rucio (ATLAS)

Community

- ATLAS
  - Approaching 400PB
  - 10M containers, 20M datasets, 1B files
  - 5K accounts, 10K identities
  - 1-2PB transfers / day, 3PB deletions / day
  - 130 sites, 600 storage endpoints
- ASGC: AMS + others
  - Several million files, 10 sites
- Xenon1T
  - 5.6 PB, 100k files, 6 sites
- Under evaluation by many communities
  - CMS, SKA, OSG (LIGO, IceCube), EISCAT_3D, FNAL (DUNE), XDC

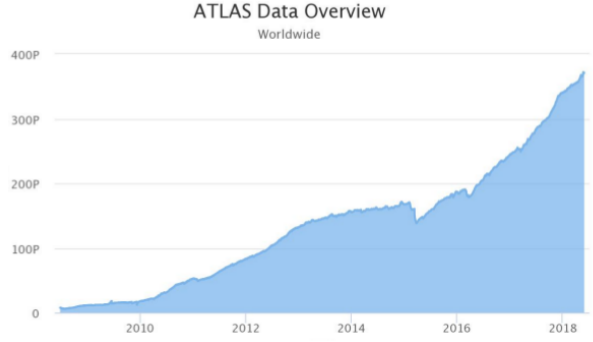

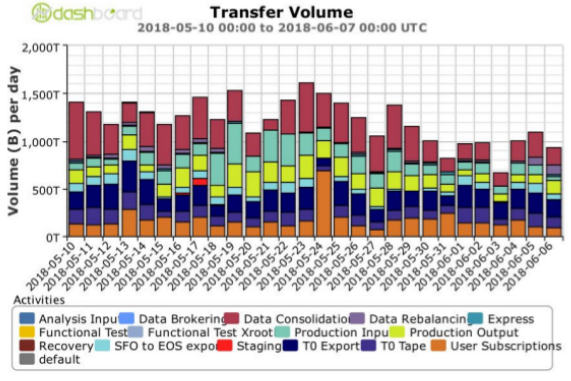

Dirac (LHCb)

DIRAC: the interware

- A software framework for distributed computing
- A **complete** solution to one (or more) user community
- Builds a layer between users and resources

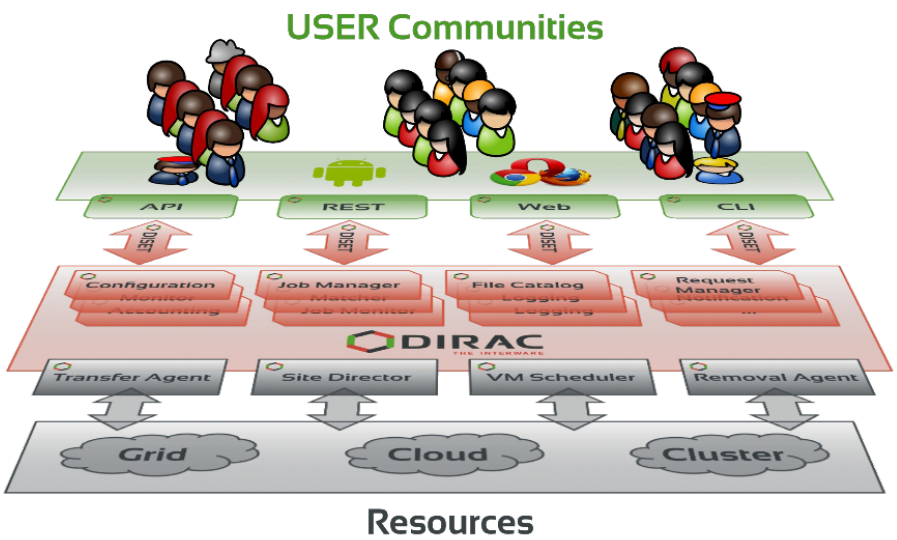

- Started as an LHCb project, experiment-agnostic in 2009
- Developed by communities, for communities
  - Open source (GPL3+), GitHub hosted, python 2.7
  - No dedicated funding for the development of the "Vanilla" project
  - Publicly documented, active assistance forum, yearly users workshops, open developers meetings
  - 4 FTE as core developers, a dozen contributing developers
- The DIRAC consortium as representing body

Dirac (LHCb)

Running the service: LHCb DIRAC Pillars

- With DIRAC, LHCb operates a service
  - need to keep a running system working, with *continuity*
- We don't see the need for a revolution
  - The system will keep evolving gradually
    - in a backward compatible way
  - Introducing new/better/faster stuff
    - Users should not notice about (most of) them

- Usability for the users
- Scalability for the services is necessary

Dirac (LHCb)

Users/communities/VOs

A *framework* shared by multiple experiments/projects, both inside HEP, astronomy, and life science

Experiment agnostic
Extensible
Flexible

Something cool

Ice Cube Neutrino Observatory

The IceCube Neutrino Observatory
Design and construction

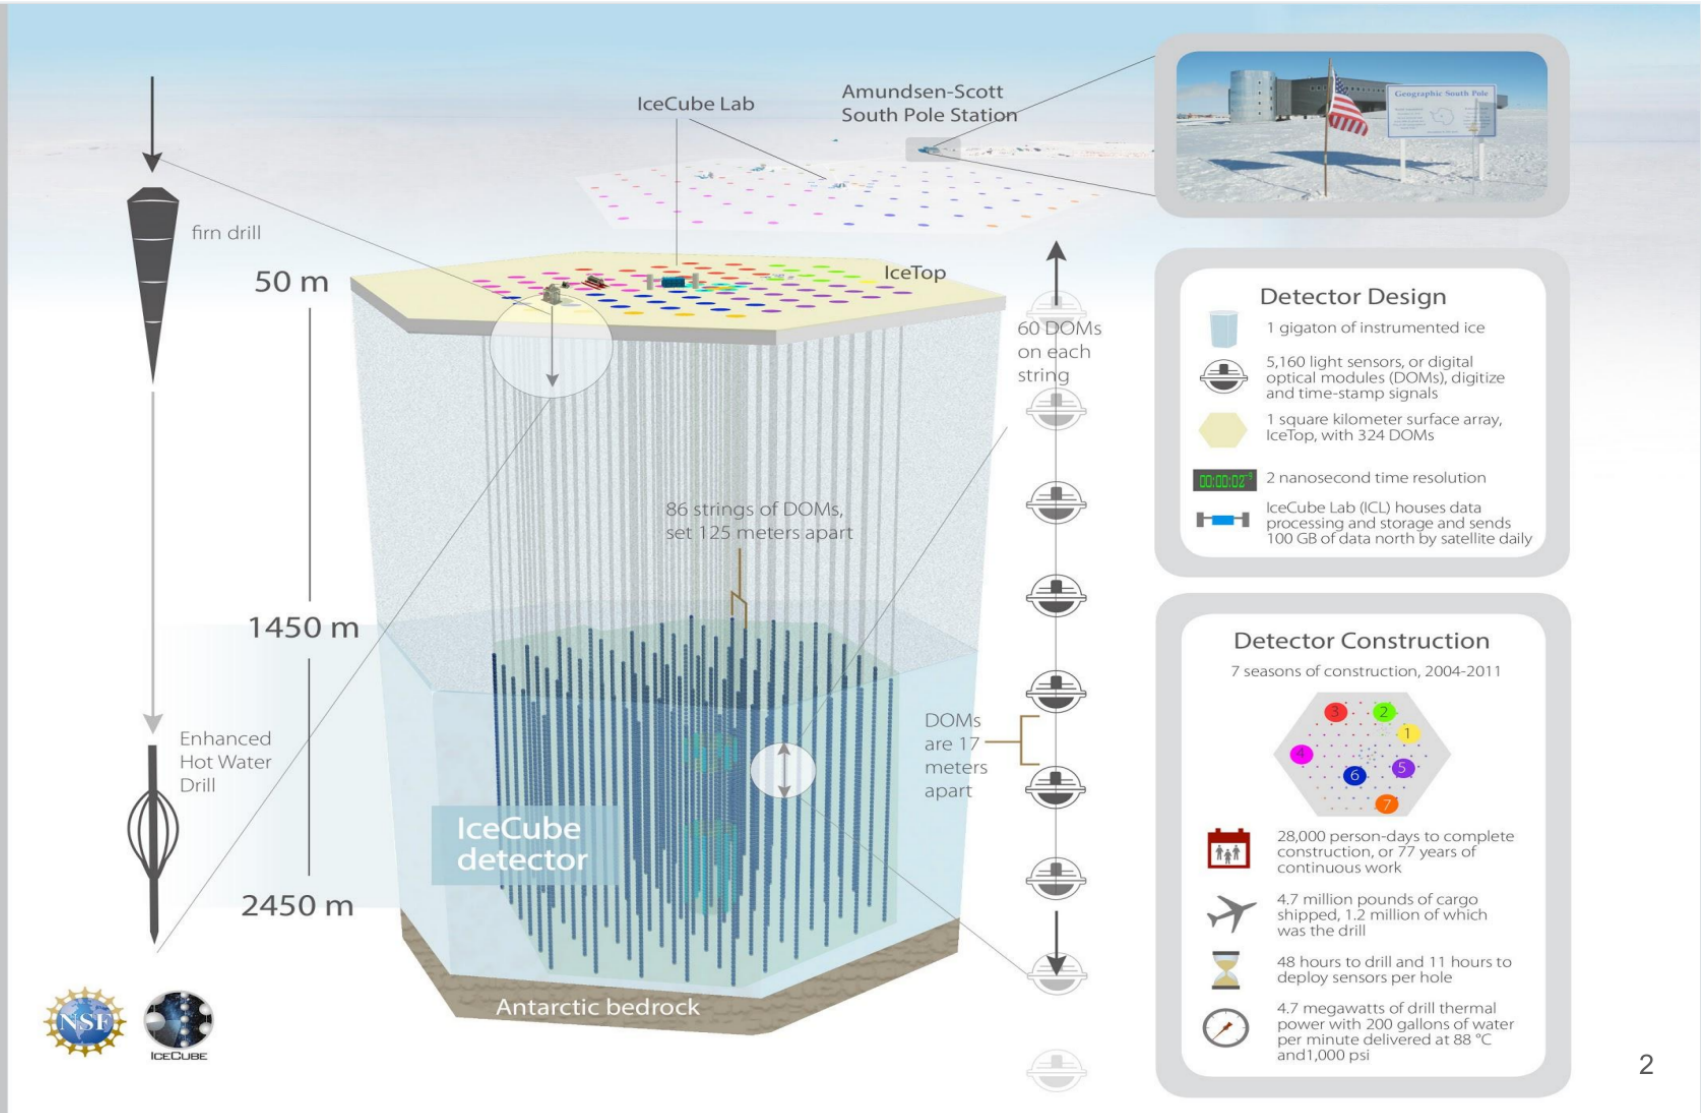

Something cool