Software and Computing Effort at BNL Physics Department

Alexei Klimentov

LSST Reps visit to BNL

Nov 7, 2019

BROOKHAVEN NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN SCIENCE ASSOCIATES

Nuclear and Particle Physics Software Group Overview

- The Nuclear and Particle Physics Software (NPPS) Group in the BNL Physics Department (~19 FTEs) collaborates and plays a leading role in software and computing development on a growing number of experiments :
  - STAR, PHENIX and sPHENIX experiments at BNL's RHIC heavy ion collider
  - The future electron-ion collider (EIC)
  - The Belle II experiment at superKEKB accelerator in Japan
  - The DUNE neutrino experiment at FNAL
  - The ATLAS experiment at the Large Hadron Collider (LHC) at CERN (Geneva, Switzerland)
- BNL has initiated in the past and currently leading projects in scientific software development at the most challenging scales
  - Workload Management System and Data Management for exascale
  - High-Performance computers and cloud computing integration with Particle Physics canonical x86 Grid centers
  - Data Carousel and automated data migration between hot/cold storage
  - Jupiter portal for physics analysis (credit to SDCC)
  - Joint R&D projects with Google and Amazon
- The group is a leader in developing solutions for the computing challenges of the High Luminosity LHC commencing in 2026, with 10 times the data rate and much greater detector and event complexity
- The group has successfully demonstrated how software and solutions initially developed for particle physics experiments can be used for other scientific intensive domains.

Outline

- Particle Physics Distributed Software stack
  - Workload and workflow management
  - Data Management
- LHC data and processing challenges
  - R&D projects
    - Data Carousel
    - Google-HEP R&D
  - High Performance Computers role for particle physics program
- Summary and conclusion

Paradigm shift in Particle Physics Computing in XXI century

Old paradigms	New ideas
• Distributed resources are **independent entities**	• Distributed resources are seamlessly **integrated worldwide** through a single submission system • **Hide middleware** while supporting diversity
• Groups of users utilize **specific resources** (whether locally or remotely)	• Access to all **resources** may be granted to all users
• Fair shares, priorities and policies are managed **locally**, for each resource	• **Global** fair share, priorities and policies allow efficient management of resources
• **Uneven user experience** at different sites, based on local support and experience	• **Automation, error handling, and other features** improve user experience • Central support coordination
• Privileged users have access to **special resources**	• All users have access to **same resources**

The Worldwide LHC Computing Grid

Tier-0
(CERN and *Hungary*):
data recording,
reconstruction and
distribution

Tier-1: permanent
storage, re-processing,
Analysis
T0 spill-over
HLT
MC Simulation
Derivation production

Tier-2: Simulation,
end-user analysis
Re-processing
Derivation production

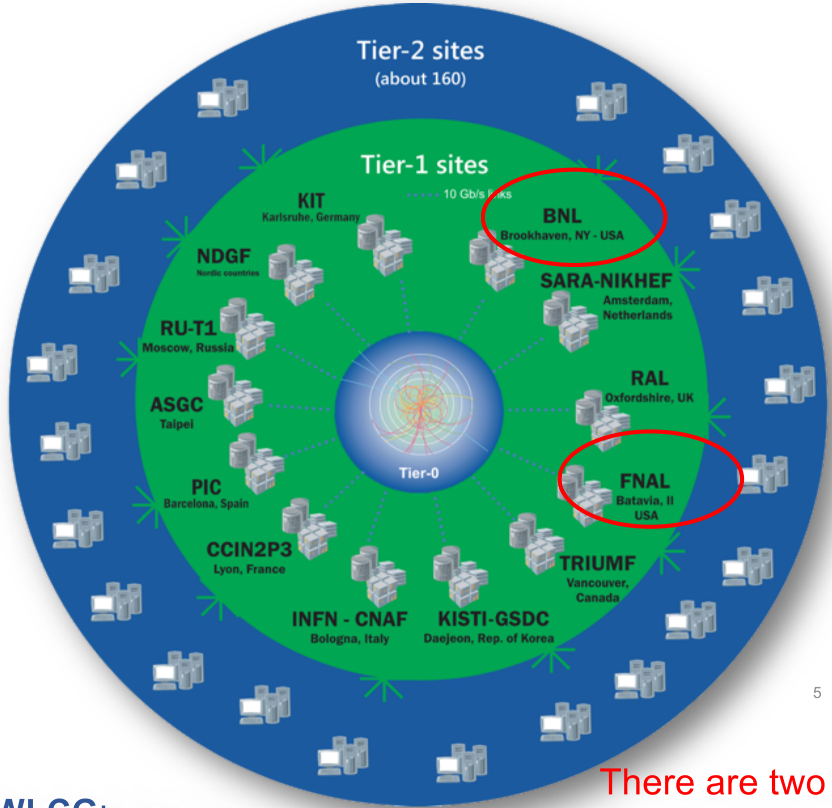

~170 sites,
42 countries

~750k CPU cores

~1 EB of storage

> 2 million jobs/day

10-100 Gb links

There are two Tier-1s in USA : FNAL and BNL

WLCG:
An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Primary distributed computing software tools

Workflow Management:
"translates" physicist requests into production tasks

STAR, (s)PHENIX,ATLAS, Belle II

Workload Management:
submission and scheduling of jobs & tasks

PanDA : AMS, ATLAS, COMPASS, NICA, nEDM…

Monitoring production jobs & tasks, shares, users

STAR monitoring, BigPanDA : AMS, ATLAS, COMPASS,…

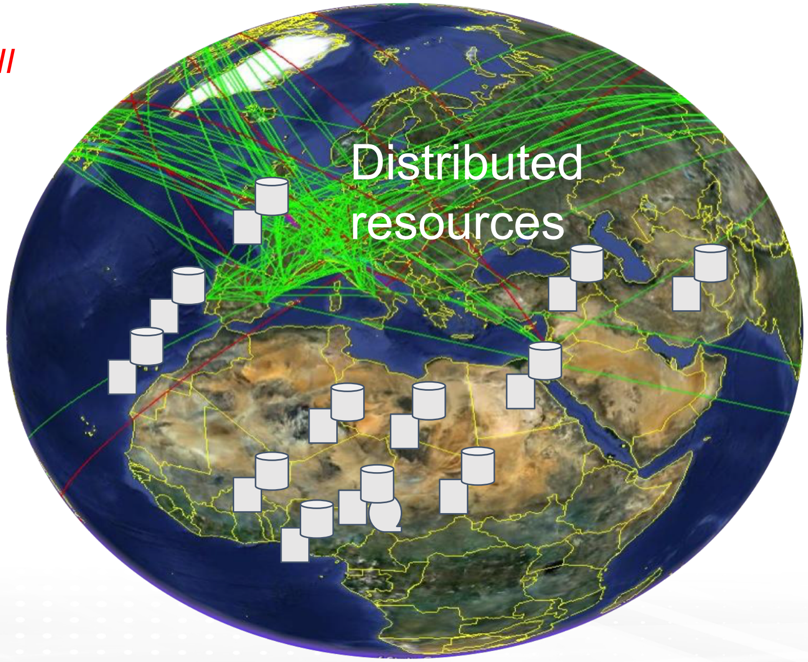

Data Management:
bookkeeping and distribution of files & datasets

Rucio : ATLAS, CMS, AMS, Belle II

Information System
(ORACLE backend)
Queues and resources description

AGIS : ATLAS, CMS

Databases: Conditions and data processing (ORACLE, mySQL, PostgresSQL)

STAR, (s)PHENIX,ATLAS, Belle II, DUNE

For six principal distributed computing software systems BNL has led development for all of them for many experiments

ATLAS Workflow and Workload Management. ProdSys2/PanDA

Global ATLAS operations
Up to ~1.2M concurrent jobs
25-30M jobs/month at >250 sites
~1400 ATLAS users

Orchestrate all ATLAS Workflows :

- *MC Production*
- *Physics Groups WF*
- *Data reprocessing*
- *T0 spill-over*
- *HLT processing*
- *SW validation*
- *User's analysis*
- *ART*

Shares/priorities

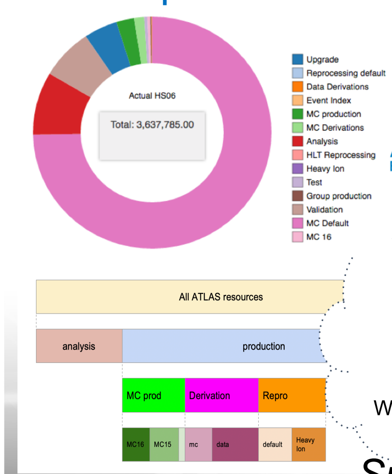

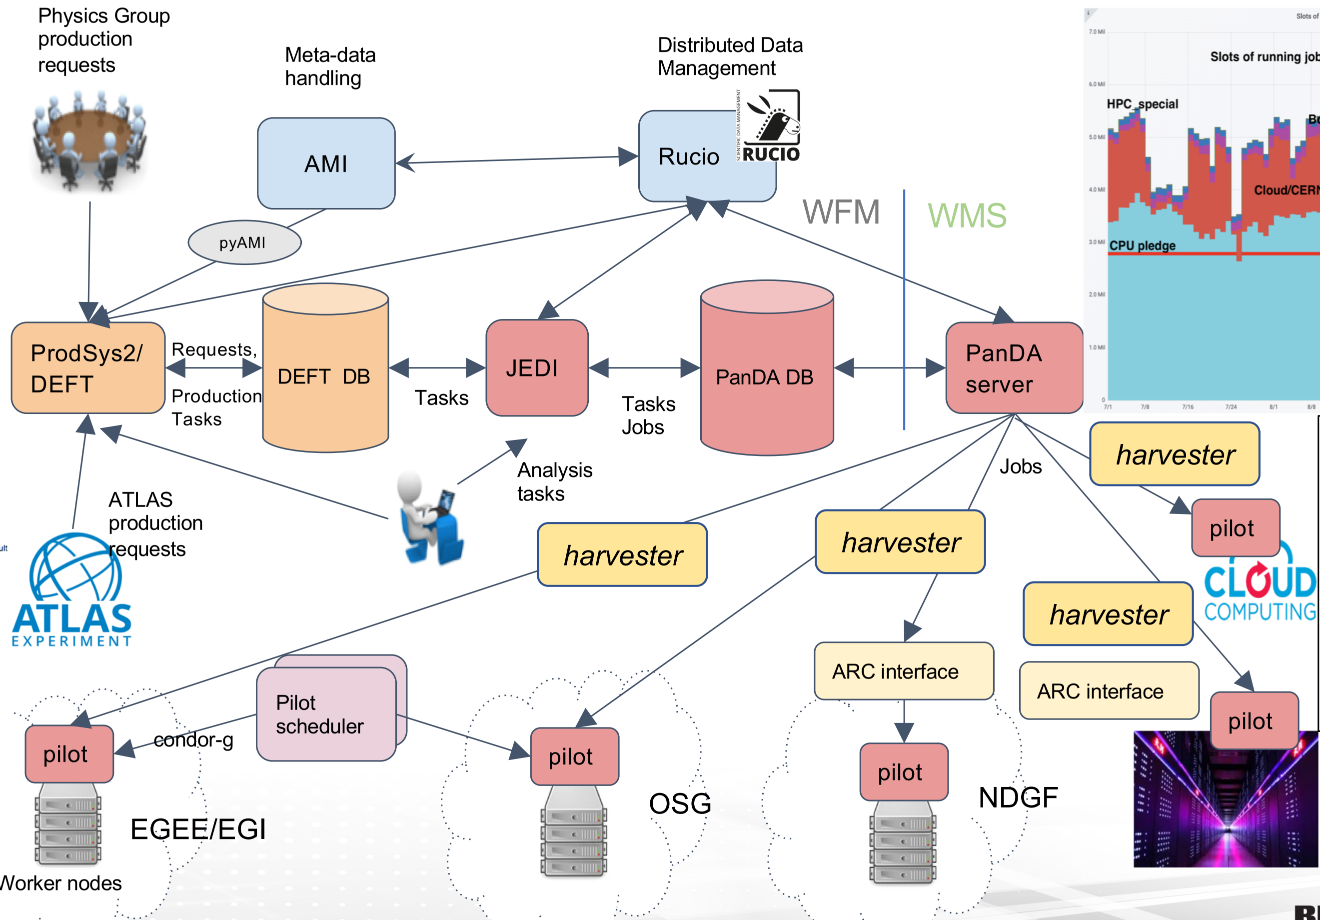

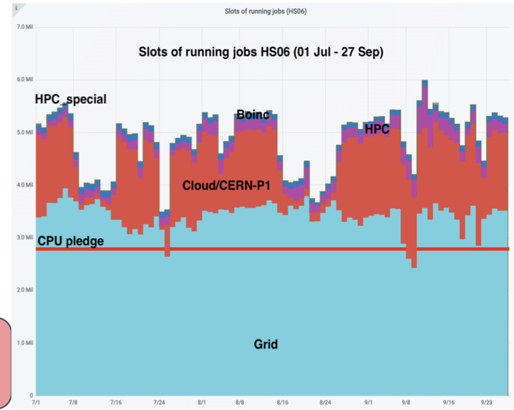

First exascale workload manager in HENP
1.4+ Exabytes processed yearly in 2014/18
Exascale scientific data processing today

Support rich harvest of heterogeneous resources. Integrate WF and data flow

BROOKHAVEN NATIONAL LABORATORY

Workload Management System Summary and Lessons learned

- ***We designed and implemented a scalable, flexible, automated production that follows physics priorities***
  - Steady state production 24x7x365 with ~300-350k cores across ~140 sites
  - HPC peaks to >1M cores, demonstrating extreme scalability of PanDA
  - The system orchestrates ~10 principal workflows and dozens of variants, with automated shares that follow ATLAS physics priorities and allocate work across global resources
  - Also supporting over 1000 analysis users with fair sharing of resources
- ***Integrated workflow and dataflow***
  - Moving >1 PB, >20 GB/s, 1.5-2M files per day
  - 405PB disk+tape, 1+B files in total (and ~540PB in 2019)
  - **PanDA processes over 1.5 Exabytes per year**
- WMS is designed by and serves the physics community
- WMS new features are driven by experiment operational needs
- WMS functionality is important as scalability
- Computing model and computing landscape in general has changed

PanDA was adopted for many HEP and astro-particle physics experiments

There are several systems with very well defined roles which are integrated for physics computing : Information system (AGIS), DDM (Rucio), WMS (ProdSys2/PanDA), meta-data (AMI), and middleware (HTCondor, Globus…). We managed to have a good integration of all of them in PanDA.

ATLAS Data management. Rucio

- A few numbers to set the scale
  - 1B+ files, 460+ PB of data, 400+ Hz interaction
  - 120 data centres, 5 HPCs, 2 clouds, 1000 users
  - 500 Petabytes/year transferred & deleted
  - 2.5 Exabytes/year downloaded & uploaded

Rucio is evaluating or already in use for many experiments including Belle II, CMS , SKA, AMS

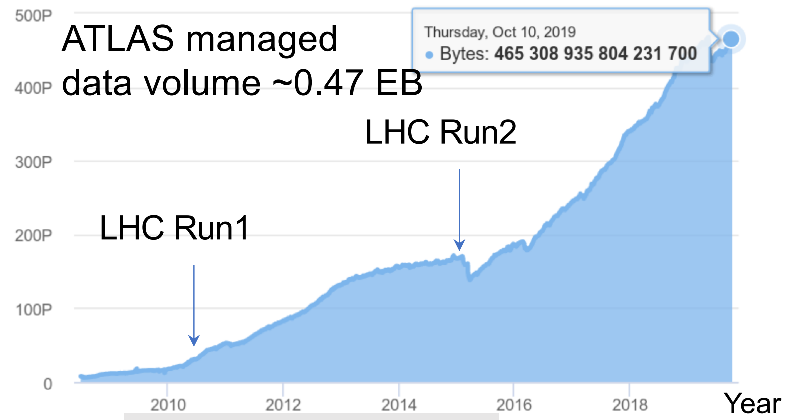

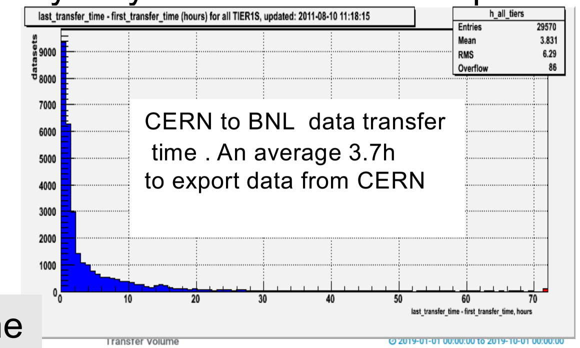

Data access volume

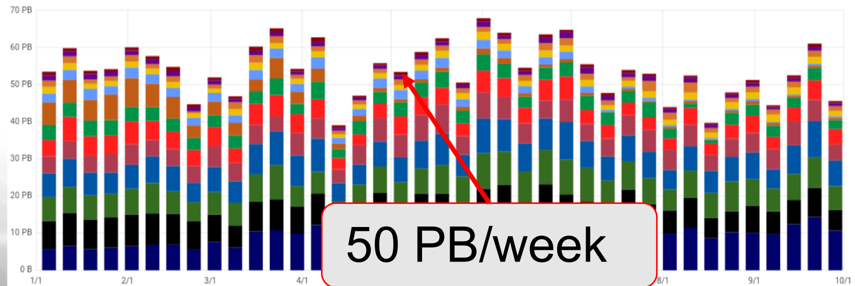

Data transfer volume

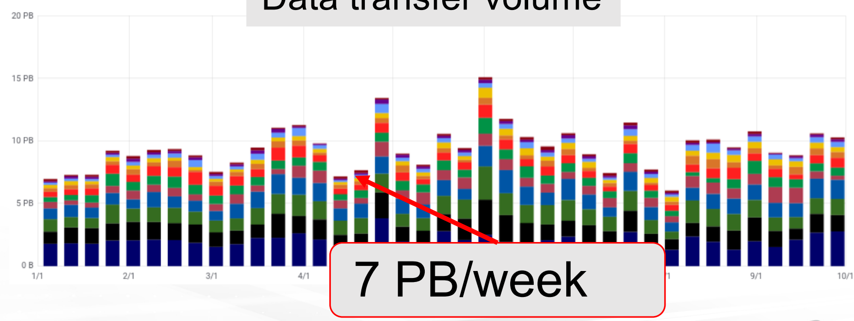

First exascale scientific data management system today

ENERGY

BROOKHAVEN NATIONAL LABORATORY

Monitoring and Analytics (bigpanda.cern.ch)

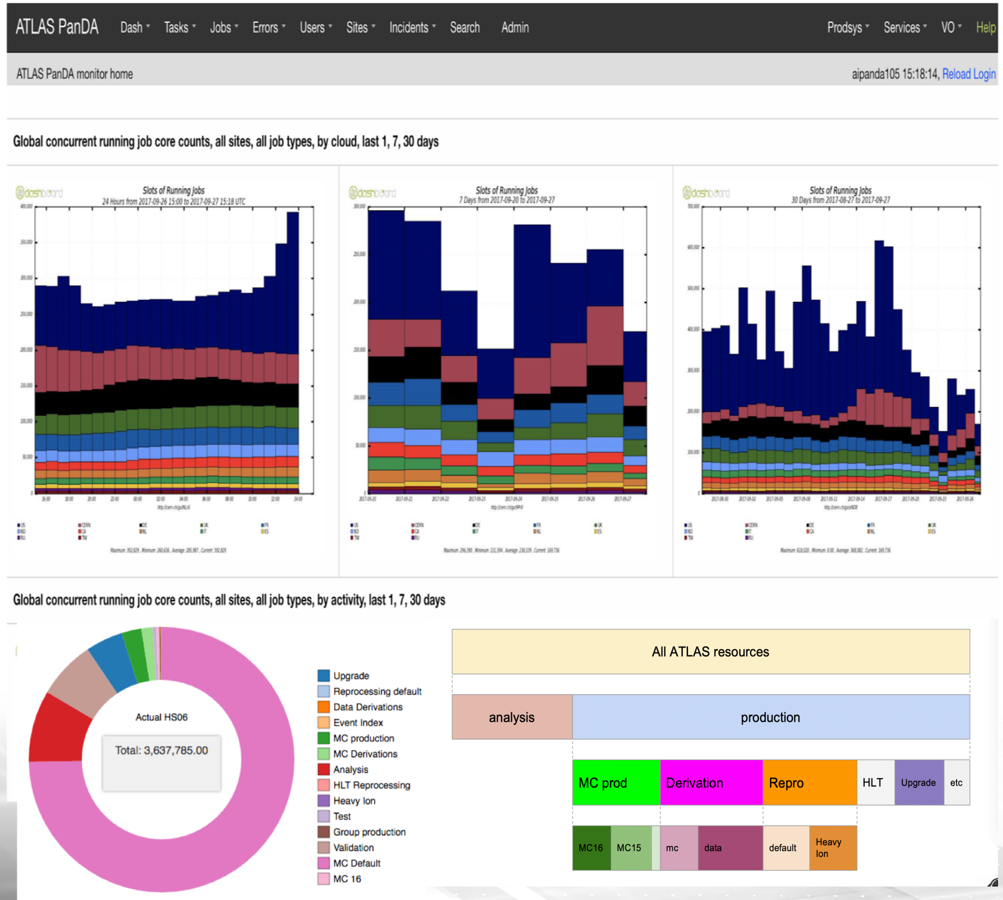

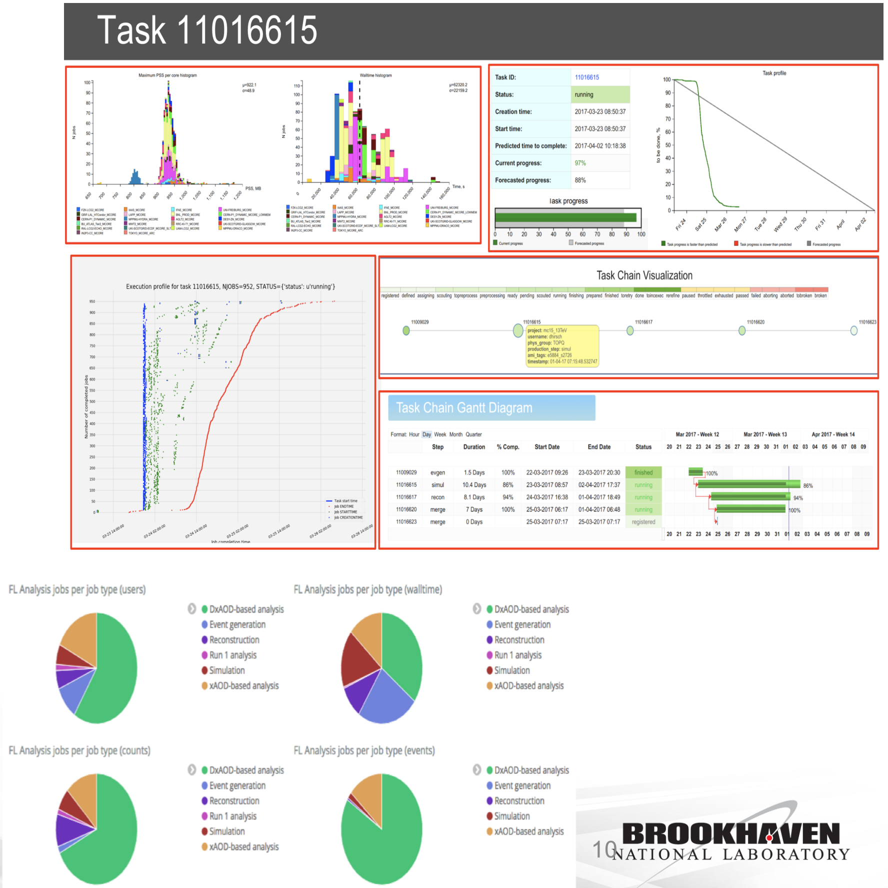

BROOKHAVEN NATIONAL LABORATORY

The High Luminosity LHC Challenge. R&D Projects

Growth in CPU Needed

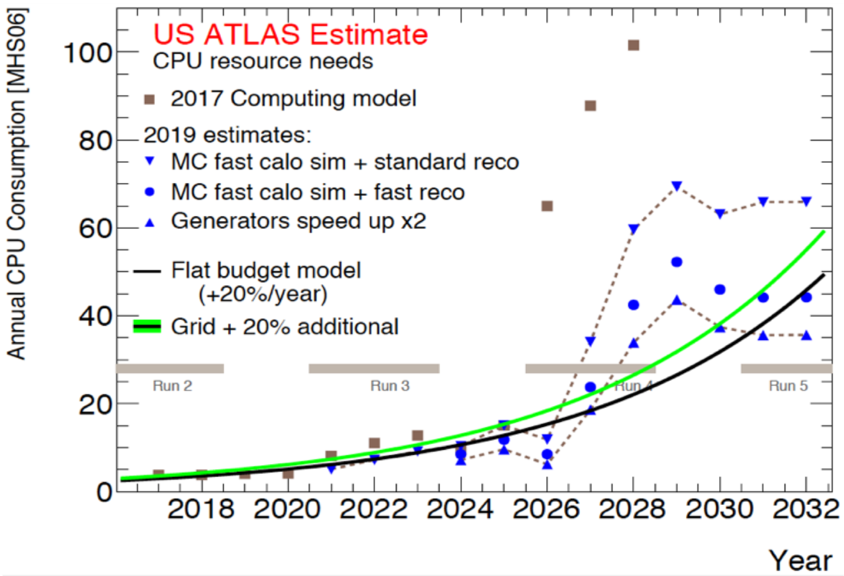

Growth in Disk Storage Needed

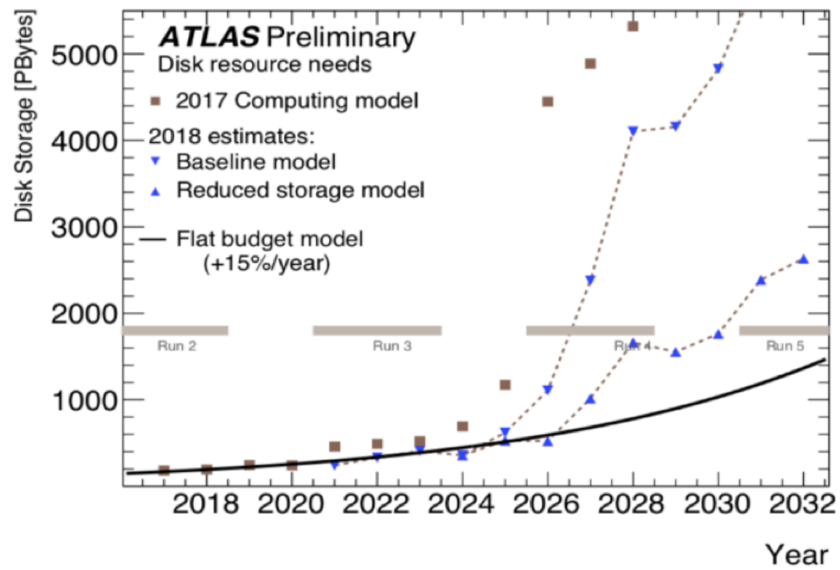

High Luminosity LHC will be a multi-exabyte challenge where the envisaged Storage and Compute needs are a factor 10 to 100 above the expected technology evolution. To address this challenge we started several R&D projects

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN NATIONAL LABORATORY

Data Carousel R&D Project

By ‘data carousel’, we mean an orchestration between workflow/workload management (WFMS), data management (DDM) and data archiving services whereby a bulk production campaign with its inputs resident on tape, is executed by staging and promptly processing a sliding window of X% (5%?, 10%?) of inputs onto buffer disk, such that only ~ X% of inputs are pinned on disk at any one time.

The project was initiated at BNL for RHIC experiments (in full production for STAR and PHENIX for more than 5 years. We manage to fetch files pretty much at tape speed for weeks in a row), today it is one of vital R&D projects to address High Luminosity LHC data processing challenge

- Ultimate goal : use tape more efficient and to have it in full production for LHC Run3 (2021)
  - Cycle through tape data, processing all queued jobs requiring currently staged data
    - ‘Carousel engine’: job queue regulating tape staging for efficient data matching to jobs?
    - Brokerage must be globally aware of all jobs hitting tape to aggregate those using staged data
  - No pre-set target on tape throughput, instead, we focus on ***efficiently*** using the ***available*** tape capacities
    - Introduce no or little performance penalty to tape throughput, after integrating tapes into our workflow
    - Improve efficiency and throughput of tape systems, by orchestrating the various components in the whole system stack, starting from better organization of writing to tapes
    - Solutions should scale proportionally with future growth of capacities of tape resources
- ‘Data Carousel’ LHC R&D was started in the second half of 2018 → to study the feasibility to use tape as the input to various I/O intensive workflows, such as derivation production and RAW data re-processing

…and “tape” could be any “cold” storage (it is led by BNL (A.Klimentov, X.Zhao; SDCC and NPPS)

HEP-Google R&D. Motivation

- IT landscape has changed dramatically since end of XX century
- US technology sector is recognized as world leaders
  - Amazon, Google, Microsoft, Oracle, … - already play significant role in worldwide scientific computing
- HENP data intensive computing challenges are (and have been) at the cutting edge of technology development
- Foster partnerships with US industries in research and development – and not just as late stage product adopters
- The huge challenges at the HL-LHC have spurred new efforts in US ATLAS to collaborate with technology partners
- Traditionally, US ATLAS Ops program did not support R&D with private sector – we are starting a new front in LHC R&D, with companies willing to invest in open source solutions

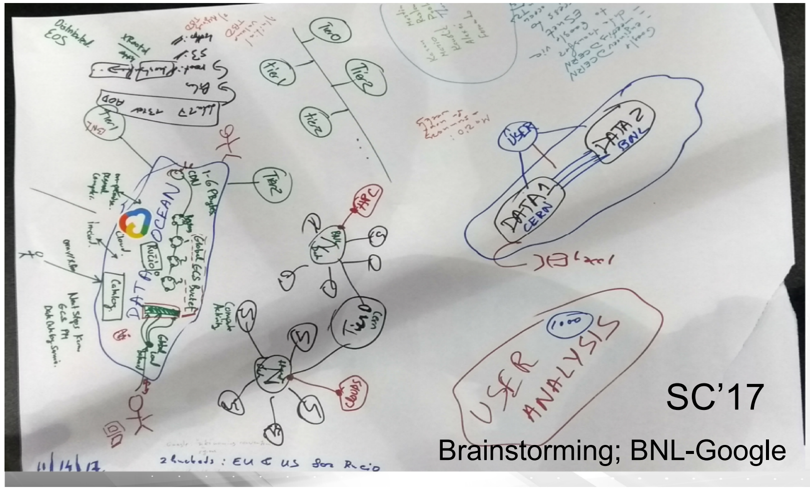

SC'17

Brainstorming; BNL-Google

Track 1	Data Management across Hot/Cold storage
Track 2	Machine learning and quantum computing
Track 3	Optimized I/O and data formats
Track 4	Worldwide distributed analysis
Track 5	Elastic computing for WLCG facilities

2018 PoC and the first realistic demo, 2019 5 Research tracks, BNL leads and coordinates three of them

High Performance Computing

- *Get extra computing resources on demand*
- *The Worldwide LHC Computing Grid and a leadership computing facility (LCF) are of comparable compute capacity.*
  - *WLCG (ATLAS share): 300,000's x86 compute cores*
  - *Titan: 300,000 x86 compute cores and 18,000 GPUs*

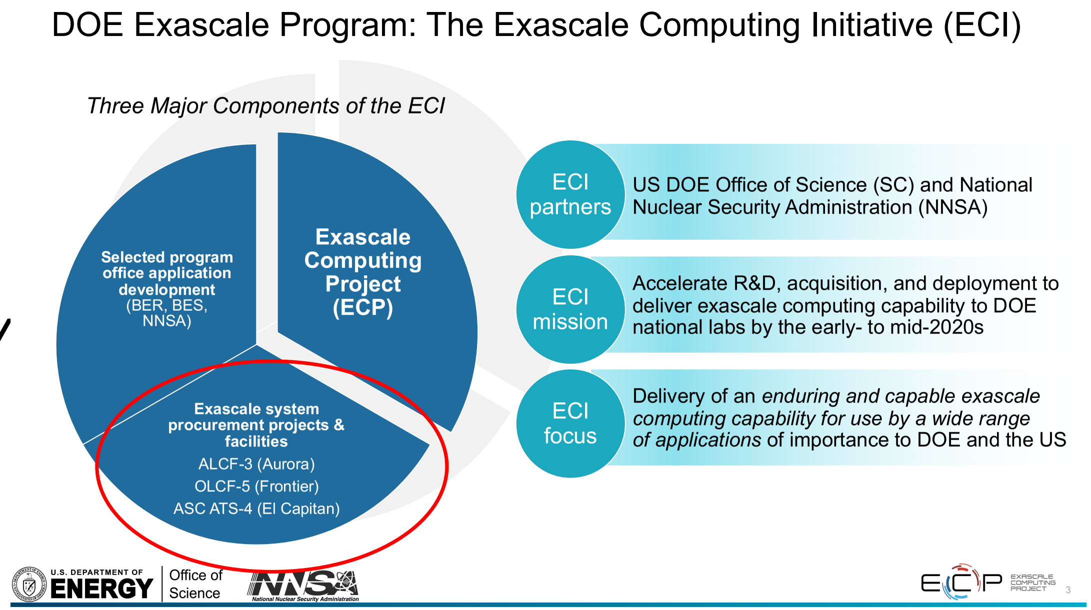

How do we efficiently integrate HPC resources and run canonical physics workflows on them ?

HPC Impact on ATLAS Production

Calendar Year 2018 Number of events Processed

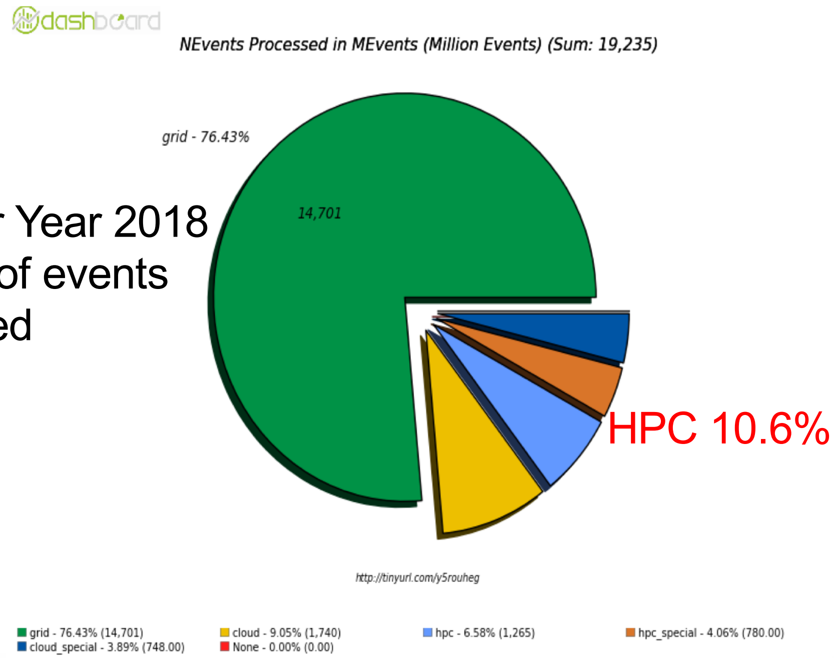

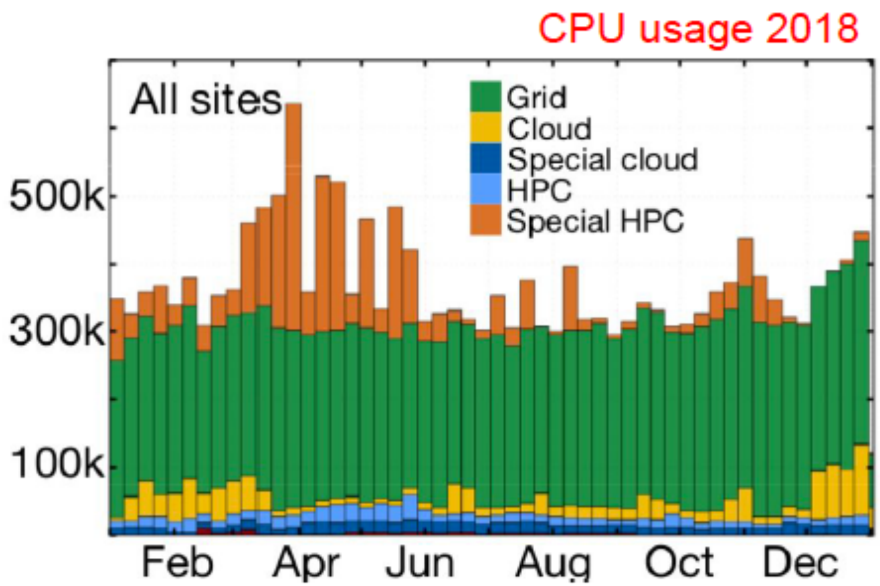

HPC facilities are also used by STAR experiment (one of the first realistic use-cases in NP)

BigPanDA Workflow Management on OLCF for High Energy and Nuclear Physics and for Future Extreme Scale Scientific Application. BigPanDA project (2012-2019) a DOE ASCR and HEP funded project since 2012; a collaboration between BNL, UTA, ORNL and Rutgers University.

USQCD

Quantum chromodynamics (QCD) is the component of the Standard Model of elementary particle physics that governs the strong interactions. It describes how quarks and gluons, the fundamental entities of strongly interacting matter, are bound together to form strongly interacting particles, such as protons and neutrons, and it determines how these particles in turn interact to form atomic nuclei.

nEDM

The goal of the nEDM experiment at the Fundamental Neutron Physics Beamline at the Spallation Neutron Source (ORNL) is to further improve the precision measurement of neutron properties by a factor of 100 to search for violations of fundamental symmetries and to make critical tests of the validity of the Standard Model of electroweak interactions

LSST DESC

The goal of the Large Synoptic Survey Telescope project is to conduct a 10-year survey of the sky that will address some of the most pressing questions about the structure and evolution of the universe and the objects in it:

- Understanding Dark Matter and Dark Energy
- Hazardous Asteroids and the Remote Solar System
- The Transient Optical Sky
- The Formation and Structure of the Milky Way

PMI

Molecular Dynamics: simulations of enzyme catalysis, conformational change, and ligand binding/release in collaboration with research group from University of Texas at Arllington.

cbi – The Center for Bioenergy Innovation

In collaboration with Center for Bioenergy Innovation at ORNL, the PanDA based workflow for epistasis research was established. Epistasis is the phenomenon where the effect of one gene is dependent on the presence of one or more modifier genes.

IceCube – South Pole Neutrino Observatory

IceCube collaborators address several big questions in physics, like the nature of dark matter and the properties of the neutrino itself. IceCube also observes cosmic rays that interact with the Earth's atmosphere, which have revealed fascinating structures that are not presently understood.

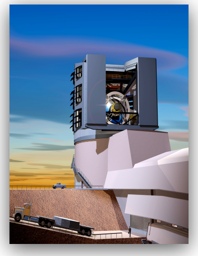

LSST Dark Energy Science Collaboration

- Collaboration with LSST/DESC since 2013 in BigPanDA project context
- The LSST Science Pipelines can process data from several telescopes using LSST's algorithms
- Pipeline to PanDA WMS submission has been implemented and tested with a standard DESC simulation workflow

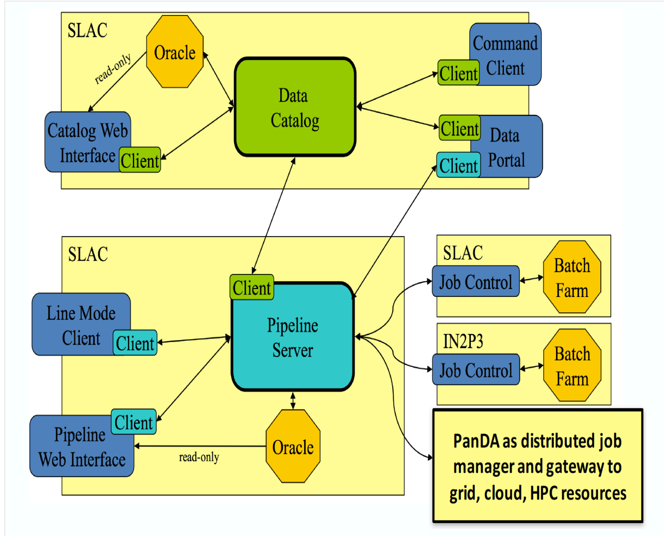

Early PoC prototype for SLAC/PanDA pipeline

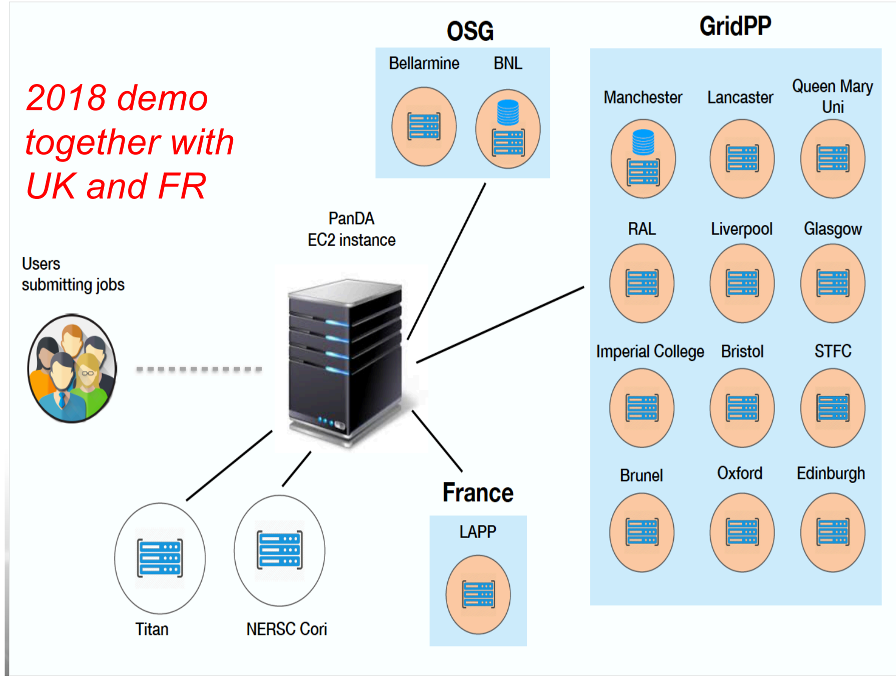

- **Facilities**
  - OSG (BNL and Bellearmine)
  - GridPP
    - 31 endpoints on 12 sites
  - LAPP Annecy (France)
  - LSST/DESC storage at 7 EU sites
    - Production results are transferred automatically to NERSC and BNL
- Evaluate Pegasus/PanDA integration

SW&Computing effort at BNL

- Core expertise in offline software and databases
  - core software expertise including its multiprocessing and multithreading variants
  - deep expertise on the C++ architecture of core HENP software and C++ itself
  - deep expertise in databases, ROOT I/O, monitoring
- Leading roles in HENP distributed software in STAR, (s)PHENIX, ATLAS and computing since its inception
  - (s)PHENIX taxi analysis platform, it enables users to run analysis with absolutely minimal maintenance needs
  - PanDA workload management system manages all distributed production and analysis workflows
  - Many innovations to grow the resources available to HENP (HPCs, clouds, fine grained processing). Pioneering work to integrate HPCs and cloud computing for HENP experiments
- Distributed software tools and systems development : PanDA, Rucio, AGIS,.. and their integration with HTCondor, Globus, Dirac, Pegasus
- Software infrastructure support (software release building and management), conditions databases
- Flexible and rapid adaptation of new software products (Belle II software stack migration to BNL was done in several months; PanDA/Pegasus integration)
- R&Ds to address new computing challenges in evolving IT landscape

For many years we are successfully collaborating with many communities to find the best technical and technological solutions and to conduct joint R&D projects

Back up slides

Data Carousel Project Phases

- Phase I : Tape Sites Evaluation
  - Conduct tape staging tests, understand tape system performance at sites and define primary metrics

done

- Phase II : ProdSys2/Rucio/Facilities integration
  - Address issues found in Phase I
  - Deeper integration between workflow, workload and data management systems (ProdSys2/PanDA/Rucio), plus facilities

In progress

- Phase III : Run production, at scale, for selected workflows
  - Address it in cold/hot storage context

We intended to conduct an iterative data carousel exercises, and to combine them with real production campaigns, to test new ideas and reveal possible bottlenecks

Goal : to have data carousel in production for LHC Run3 (2022)

Grid-HPC integration. Concrete Results

HPCs architecture, configuration and policies posed several challenges to the deployment of ATLAS distributed software components.

- The default model of PanDA pilot on the Grid was unfeasible for HPC/LCF
    - The problem is solved in **Pilot2.0**
- **Harvester**, a new interface, common across resource types, between resource and workload manager was developed, as a result ALCC utilization was increased.
- **ARC software**, our backbone for HPC integration in Europe. Many EU HPCs are integrated via ARC software
- ATLAS Workload Management and Distributed Management Systems (**PanDA** and **Rucio**) have successfully coped with increased workload and traffic after HPC integration
- ATLAS **software releases** have been successfully built on HPCs (Titan and Summit)
- **Event Service**. High HPC utilization via fine grained workflows

Concrete Results. New modes for use HPC for Data Intensive Sciences.

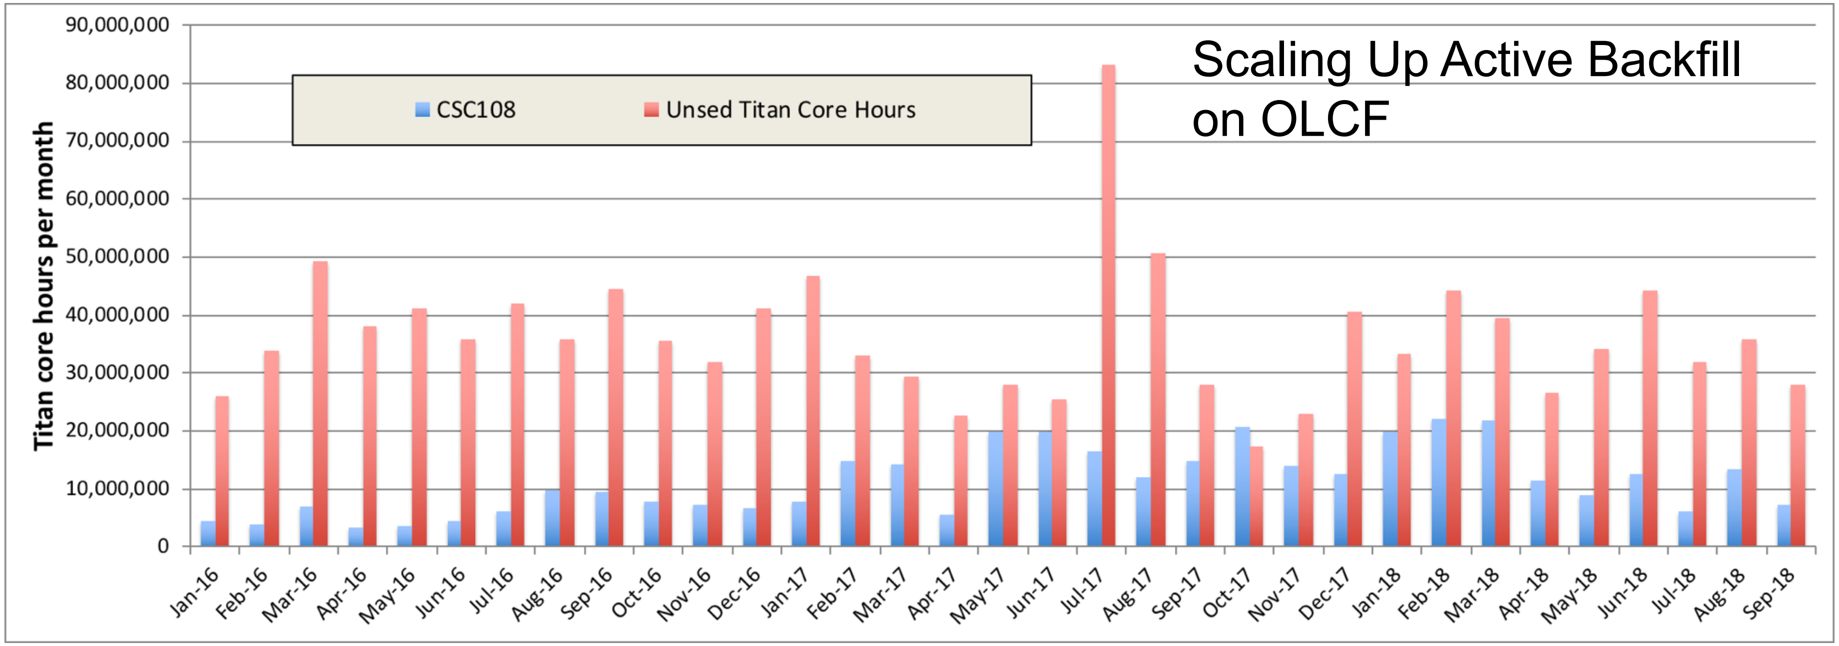

Consumed 370 Million Titan core hours from Jan. 2016 to Sep. 2018

- This is 2.8 percent of total available time on Titan over this period

Remaining used backfill slots are often too short or too small for assigned ATLAS payloads

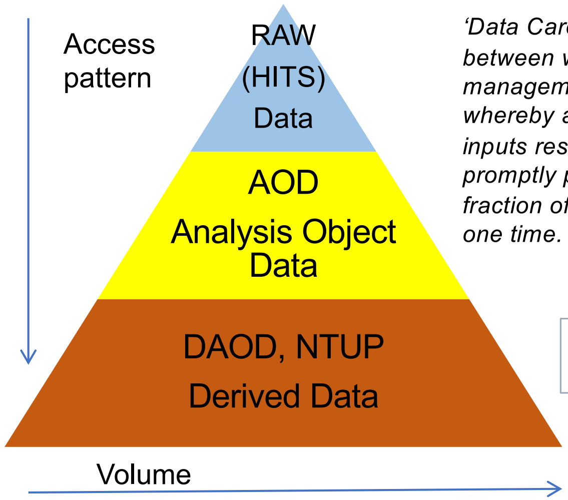

'Data Carousel' we mean an orchestration between workflow management (WFMS), data management (DDM/Rucio) and tape services whereby a bulk production campaign with its inputs resident on tape, is executing and promptly processing of inputs. Only a small fraction of inputs are pinned on disk at any one time.

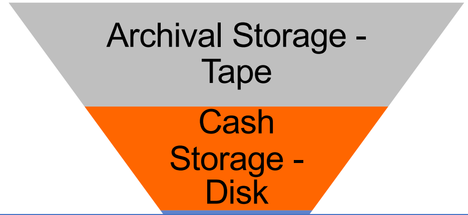

Data Carousel Model. Automatic data migration between disk and tape

Hot/Cold Storage Model

Hot/Cold storage model gives us more flexibility with data handling .
We can archive ALL data on tape and keep on disk and cash the most popular data
Plan A (α, β): Data will migrate between hot/warm/cold storage automatically
Plan B : will address the case when tape drives market will be in danger

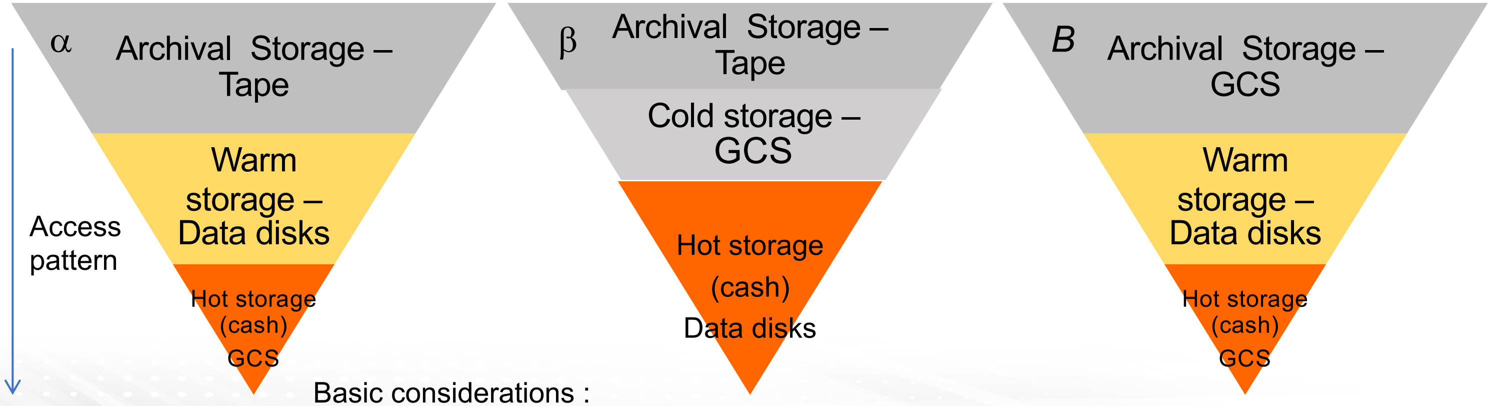

Basic considerations :

1. Access pattern
2. Cost, performance and capability
   1. Capability = functionality.. How well requirements are managed
   2. Performance = data availability, retrieval speed and data access speed

- **Five EuroHPC-JU petascale systems, 50+PFLOPS,** installed by 2020
  - Meluxina, Deucalion, Euro-IT4I, Vega, PetaSC

- **Three EuroHPC-JU pre-exascale consortia 600+PFLOPS,** installed by 2020
  - Lumi, BSC, Leonardo
  - Lumi (Kajaani Data Center, FI) : Finland, Belgium, Czech Rep, Denmark, Estonia, Norway, Poland, Sweden, Switzerland

- **Two exascale sites**