ATLAS Computing
Leveraging Heterogeneity

Kaushik De, Alexei Klimentov

CSCS, Lugano
January 10, 2018

Introduction to ATLAS

- ATLAS at CERN
  - Largest experimental detector ever built
  - Higgs discovery, precision measurements of the Standard Model (including Higgs), searching for dark matter, supersymmetry, exotic particles… >700 publications
  - LHC physics program well planned out for the next 20 years
- ATLAS Computing Challenge
  - Massive scale of computing resources required
    - Currently using about 300k batch slots and 300 PB storage continuously
  - Huge collaboration – with thousands of data analyzers
  - Complex workloads and workflows
  - Computing needs grow every year with more LHC data

The Nobel Prize in Physics 2013

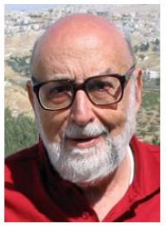

Photo: Pnicolet via Wikimedia Commons

François Englert

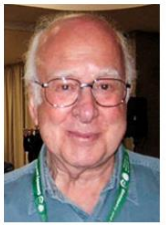

Photo: G-M Greuel via Wikimedia Commons

Peter W. Higgs

The Nobel Prize in Physics 2013 was awarded jointly to François Englert and Peter W. Higgs *"for the theoretical discovery of a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle, by the ATLAS and CMS experiments at CERN's Large Hadron Collider"*

ATLAS Computing Model

- ATLAS computing model is designed to use distributed data centers remotely from the start
  - No single facility of the required scale was available
  - Collaborators were scattered worldwide – early experience of using resources locally proved too labor intensive
- Distributed computing model is based on three main pillars
  - WorkFlow Management System - ProdSys
  - Workload (execution) Management System – PanDA
  - Data Management System – Rucio
- All three systems tightly integrated for efficient performance
- ATLAS data centers started out organized in Tiers, but advances in networking have eliminated strict hierarchy

ATLAS Computing Resources

- Computing infrastructure is built around WLCG
  - WLCG == The Worldwide LHC Computing Grid
  - Distributed High Throughput Clusters (HTC)
  - Currently ATLAS has access to about 150 such HTC clusters
  - While resources vary greatly – they are accessed using standard grid protocols (with small regional variants)
  - Cloud computing resources are also integrated into ATLAS – somewhat similar to grid sites

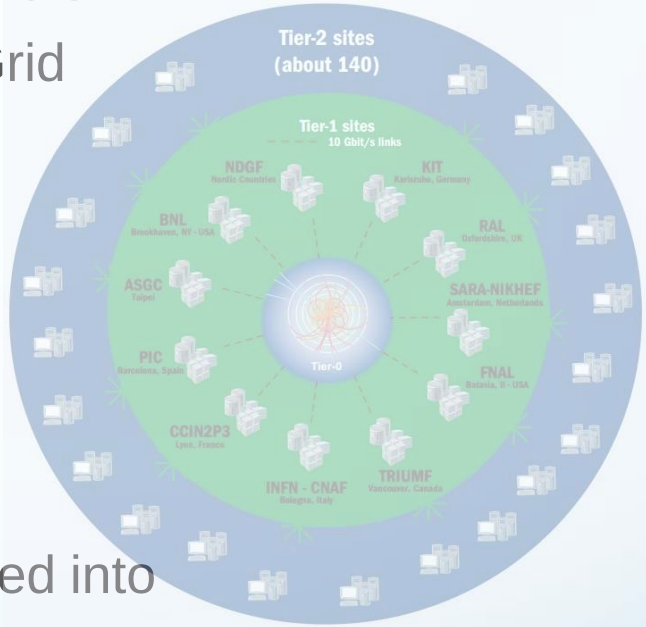

- However, even this large pool of resources is insufficient to meet ATLAS physics goals

From HTC to HPC

- With highly successful Run 2 (~x2 data delivered), and looking ahead to Runs 3, 4 at the LHC
- ATLAS started looking at traditional HPC systems
  - Almost 50% of ATLAS CPU cycles used for simulation
  - HPC architectures are well suited to run simulations
  - However, they need to be integrated into production and data management systems – not standalone
  - Next few talks will elaborate how we did it for ATLAS
- HTC/Grid + Clouds + HPC == Truly Heterogeneous and distributed computing integrated seamlessly

ATLAS CPU Consumption Since 2015

Three Tiers:

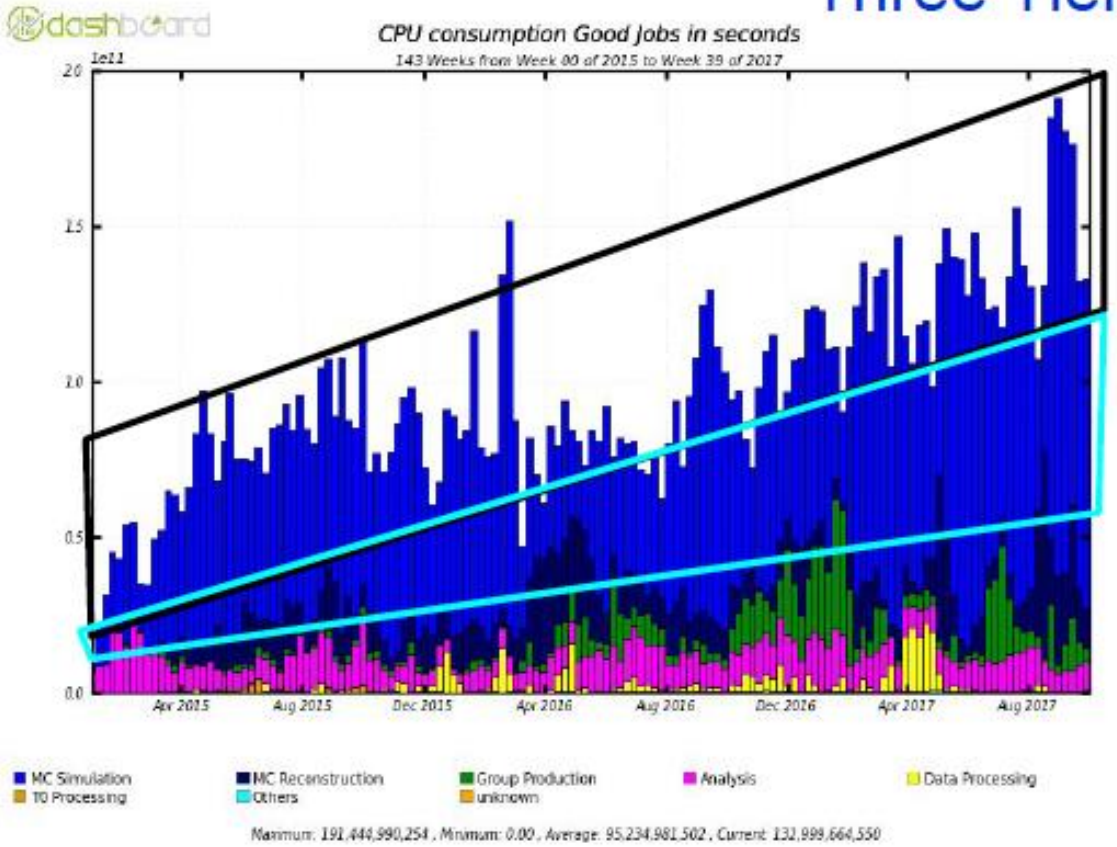

Utilization spikes. May be cost effective to offload to on-demand resources

CPU-intensive (simulation)
Well suited for HPCs and Clouds

Data Intensive Processing
Best run on grid

Why is HPC Different?

- If HPC site provides grid or cloud interface – then ATLAS can use the site easily and immediately
- Long term storage at HPC is nice – but not required
- With Harvester (described in other talks), we can enable ATLAS access through edge machine
  - This is secure and very safe route
  - Ideally need good server for edge machine
  - May require multiple edge machines for scalability
- Other issues like installation, networking, data movement… in later talks
- HPCs also require special optimizations because of their size

Benefit of HPCs for ATLAS

- Primarily, raw computing power, for simulation and reconstruction, enabling new LHC physics results
- But also new model for evolving HTC and HPC computing paradigms into unified infrastructure
- Opportunity to optimize application code
- Innovative solutions for remote data access and caching - data lake, data ocean…
- Partnership between scientific instruments & HPC
- Partnership between physicists and computer scientists

The Titan Example

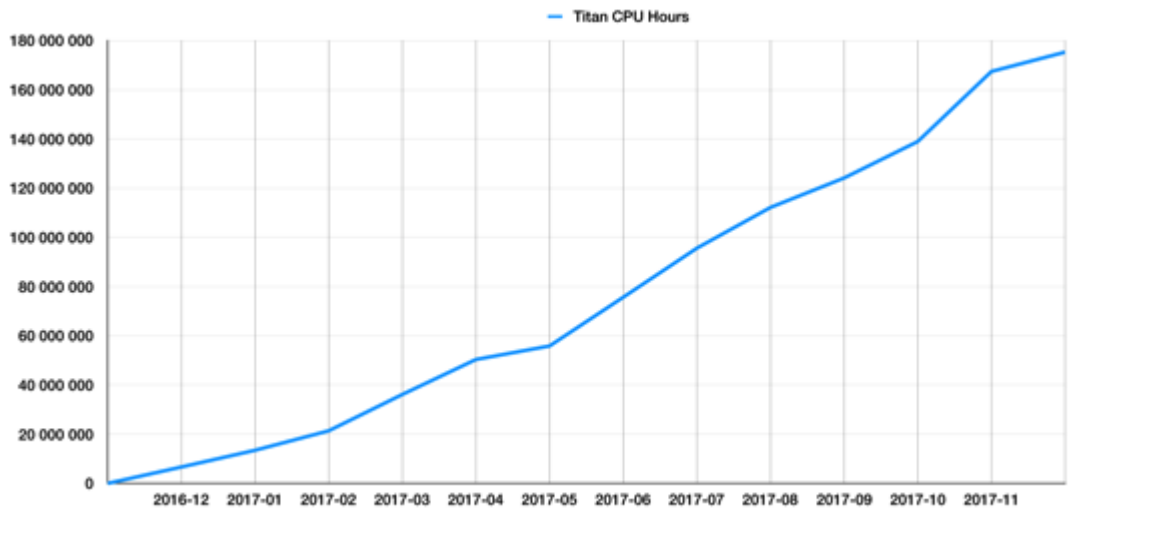

180M Titan core-hours were used by ATLAS (backfill and ALCC) In past ~12 months

Partnership between ATLAS and Titan far more valuable than core hours used

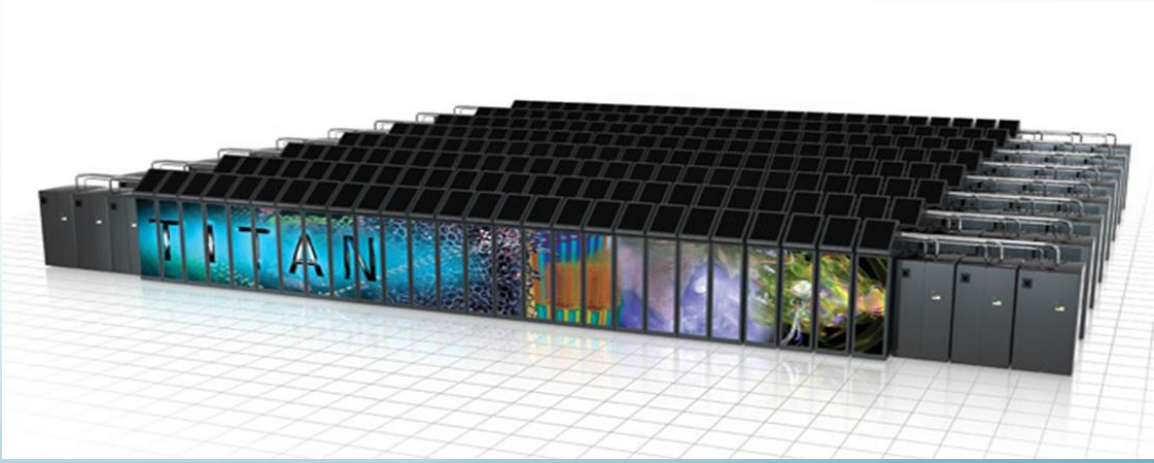

OAK RIDGE National Laboratory

Challenges Ahead

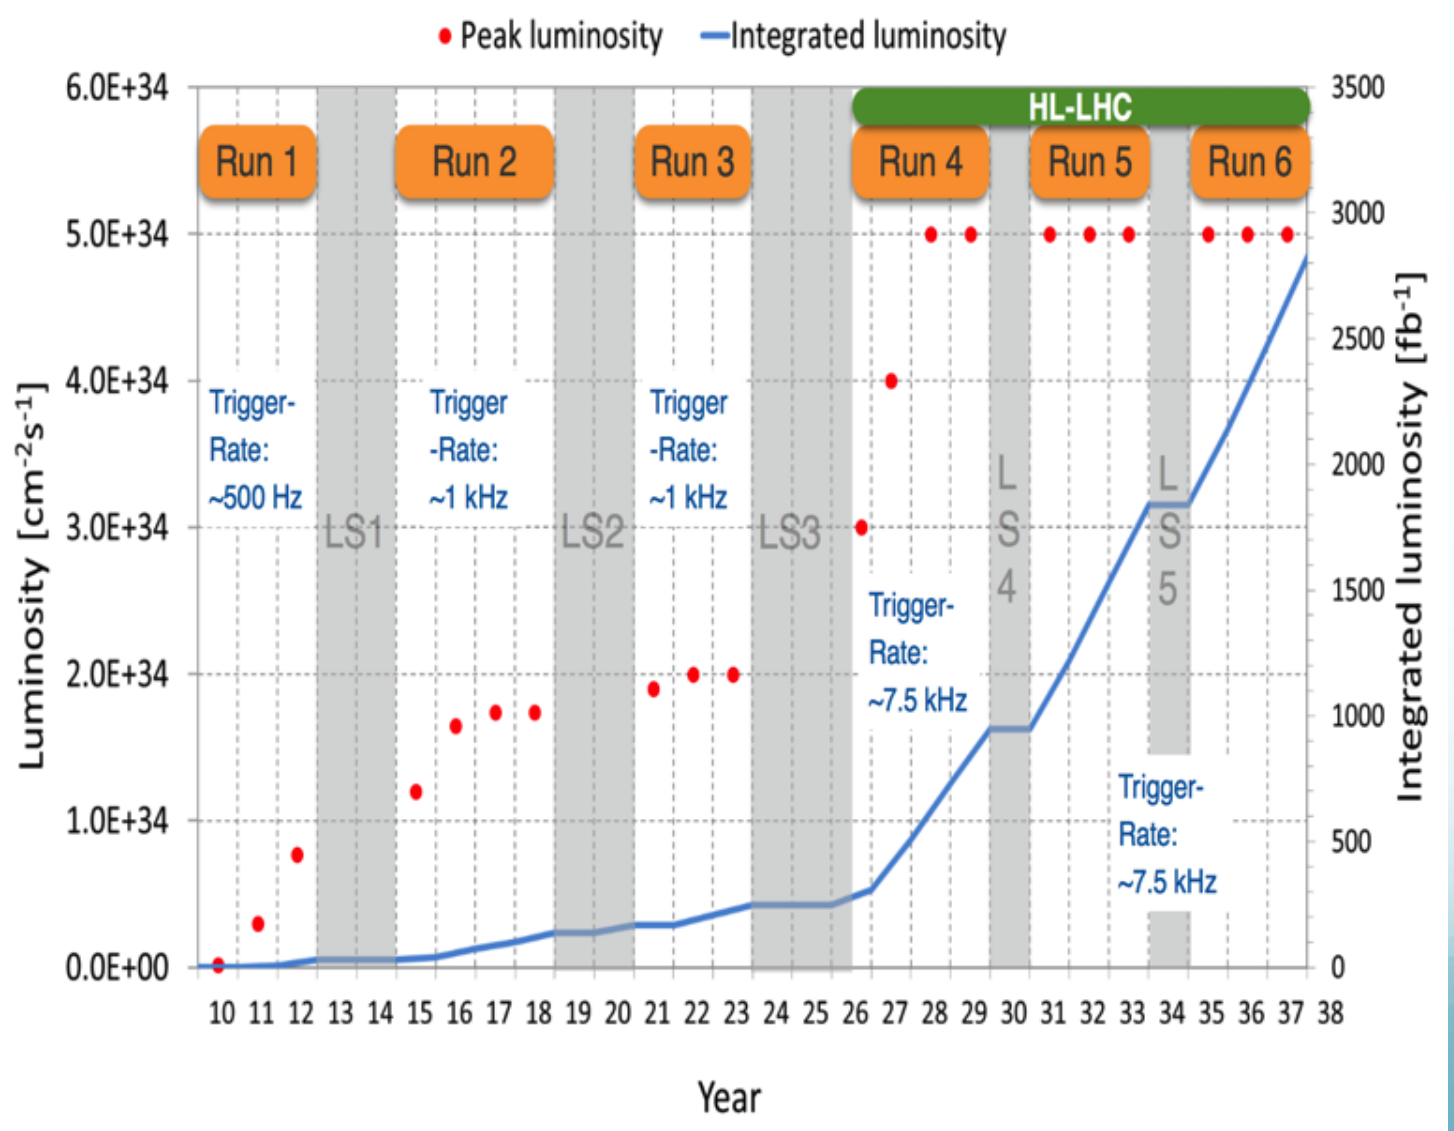

Simplistic Extrapolations