

DevOps in ATLAS

Examples from Experimental HEP

Kaushik De
University of Texas at Arlington
[ATLAS Collaboration](#)

MAGIC Meeting
August 1, 2018

Science Overview

- ATLAS experiment at CERN
 - Large Hadron Collider (LHC)
 - Largest detector ever built
 - pp collisions at 13 TeV energies
 - > 700 publications
 - Higgs discovery in 2012
 - Precision studies of the Standard Model of physics
 - Searching for dark matter, supersymmetry, exotics...
 - 25 years since inception, 20+ years more to go

The Nobel Prize in Physics 2013

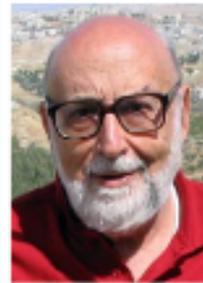


Photo: Prikolel via
Wikimedia Commons
François Englert



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*

Computing Challenge

- ATLAS computing needs
 - 1 PB data from detector every second, 1-2 GB recorded
 - 350k cores continuously used worldwide for processing, with ~1M cores peak usage
 - 200 TB disk storage, 200 TB tape storage, worldwide
 - ~100 data centers, ~1000 users
- ATLAS Software stack
 - 5M lines of physics application code, written by hundreds of people (many of them physicists), used by thousands
 - 1M lines of distributed computing code to manage automated execution, catalogue and control
 - Continuously evolving

DevOps as Evolution

- ATLAS experiment is 25 years old
 - DevOps emerged as the working model over time
 - Followed natural evolution of the experiment from concept to steady operations
 - Nurtured by the inherent culture of a large collaboration
- First 10 years, design and development focus
 - Software was meticulously planned, designed and developed into working applications
- Middle 5 years, development and operations ran side-by-side, with periodic cross interactions and iterations
- Last 10 years, as the experiment entered data collection phase, software and computing has fully embraced a DevOps model

DevOps as Requirement

- ATLAS has been steadily collecting experimental data for the past 10 years
 - Data is collected and processed with high efficiency
 - Analyzed by a thousand physicists and published quickly
 - As a frontier data science experiment, software has been constantly changing (major evolutions every few years)
 - The software and computing team is large, distributed, multi-national, and collaborative (no top down hierarchy)
- DevOps is a natural solution to large data science
 - DevOps at every level – applications, infrastructure, tools
 - Need for nightly releases, continuous testing and validation, monitoring and analytics, ticketing ...
 - Skip tools and practices here – look at examples

Some Examples

- Picked from my personal experiences
- Example 1: DevOps from the start
 - PanDA+ProdSys+Rucio – ATLAS distributed computing applications for workload and data
- Example 2: Ops working with Dev
 - Optimizing code to run at NERSC
- Example 3: Looking to the future
 - SLATE project to automate deployment and operation of distributed services

ATLAS Distributed Computing

- PanDA – Production and Distributed Analysis system
- Tightly integrated with ProdSys (task management front end) and Rucio (data management system)
- Python code, ~1M lines (all 3), agile from the start
- Core developers (dozen), operations and support team (hundred – including infrastructure teams), and users (thousands of physicists) working together
- DevOps is practiced throughout
 - Dev does Ops, Ops does Dev, all do testing and integration, shift teams, support teams – all work together in ADC

PanDA System

PanDA concepts



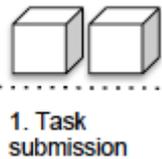
Rucio

- ▲ 2. Dataset lookup, file locations and replication requests
- ▼

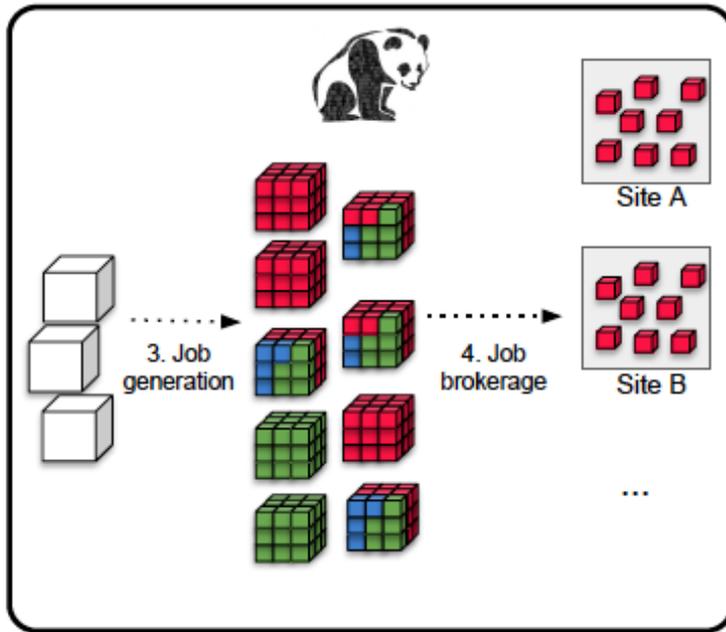
- Ready job
- Running job
- Finished job
- ✈ Pilot



Users and groups



1. Task submission

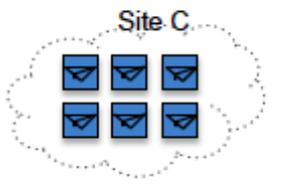
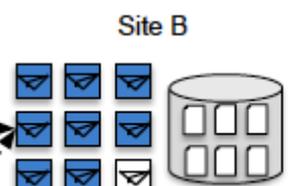
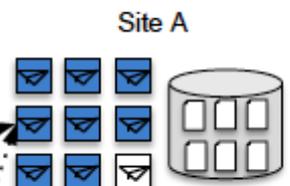
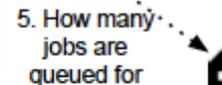
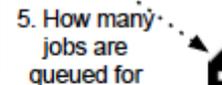
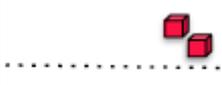


3. Job generation

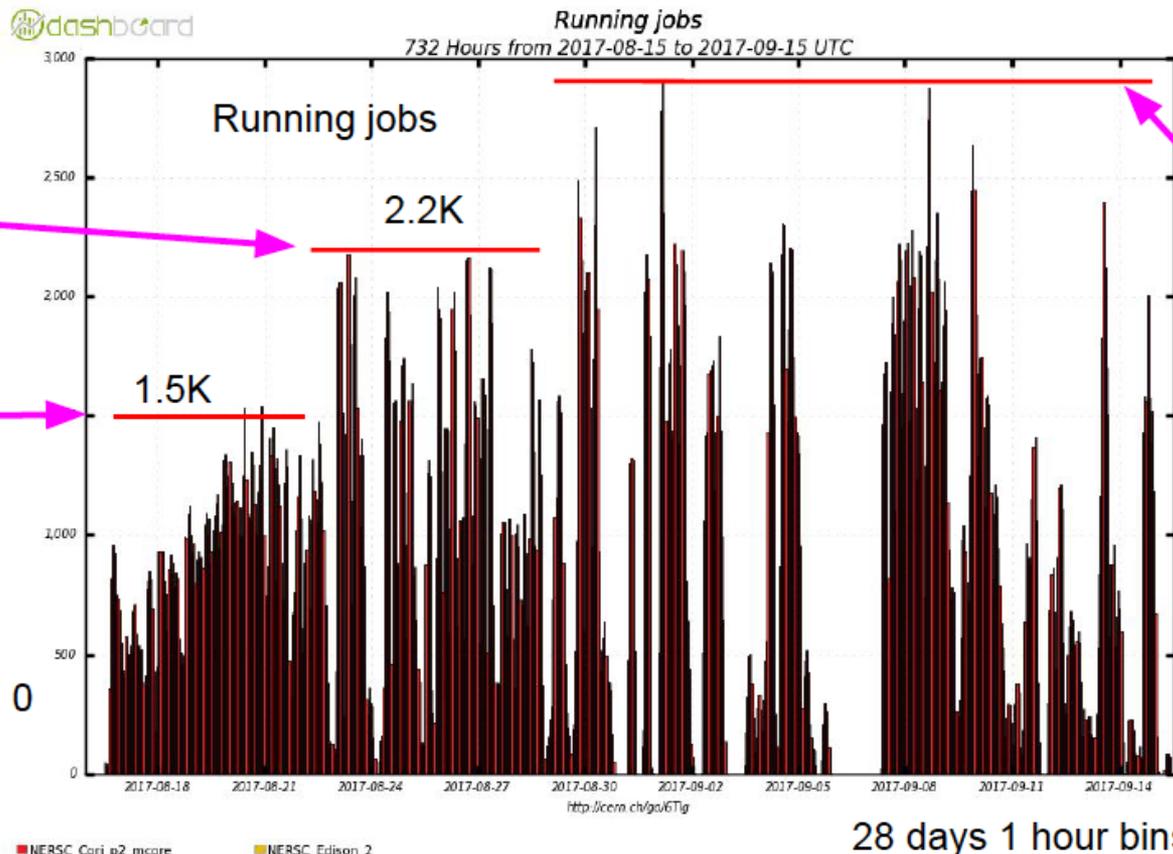
4. Job brokerage



...

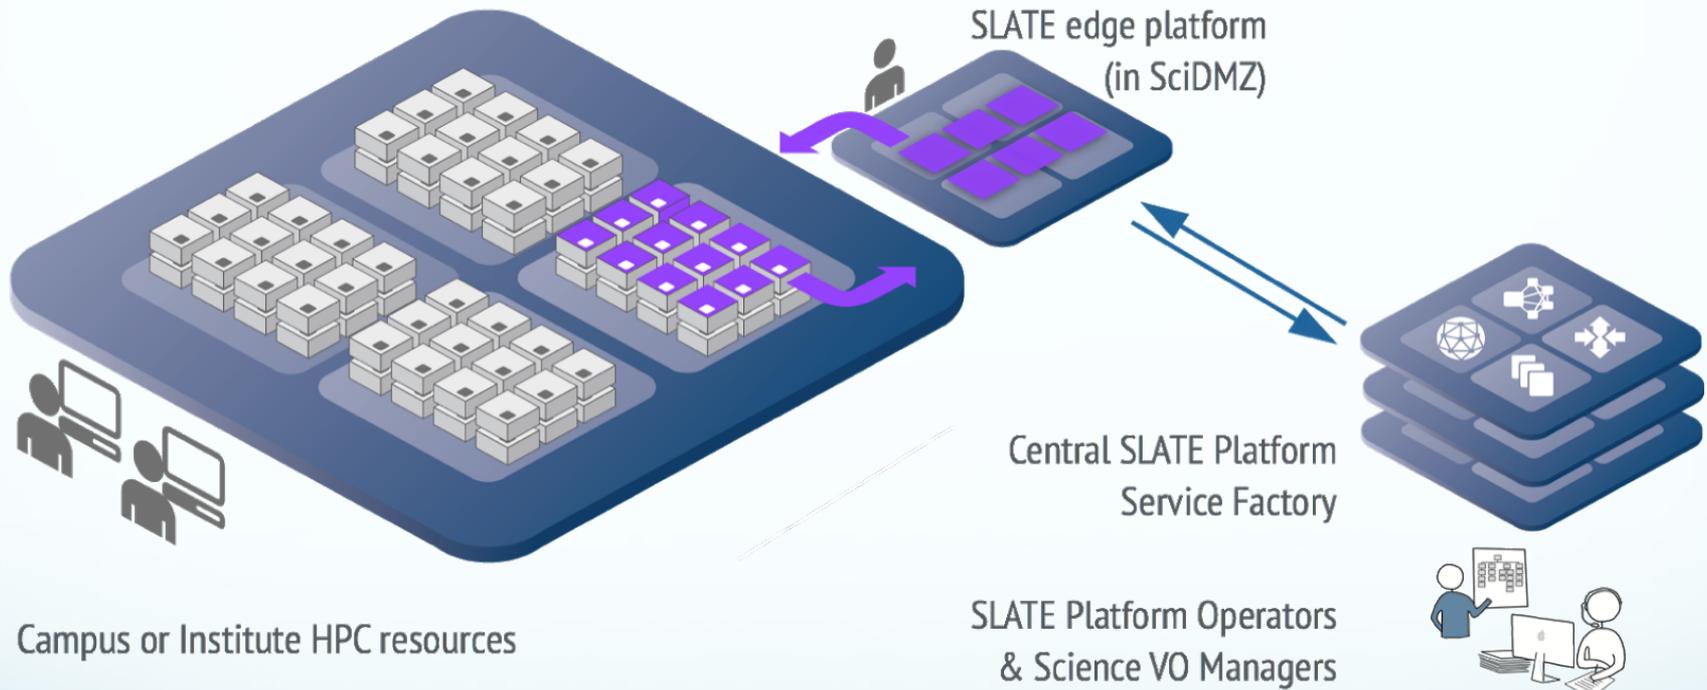


Effect of using shifter containers at NERSC



Doug Benjamin, ANL

SLATE – slateci.io



Services Layer at the Edge

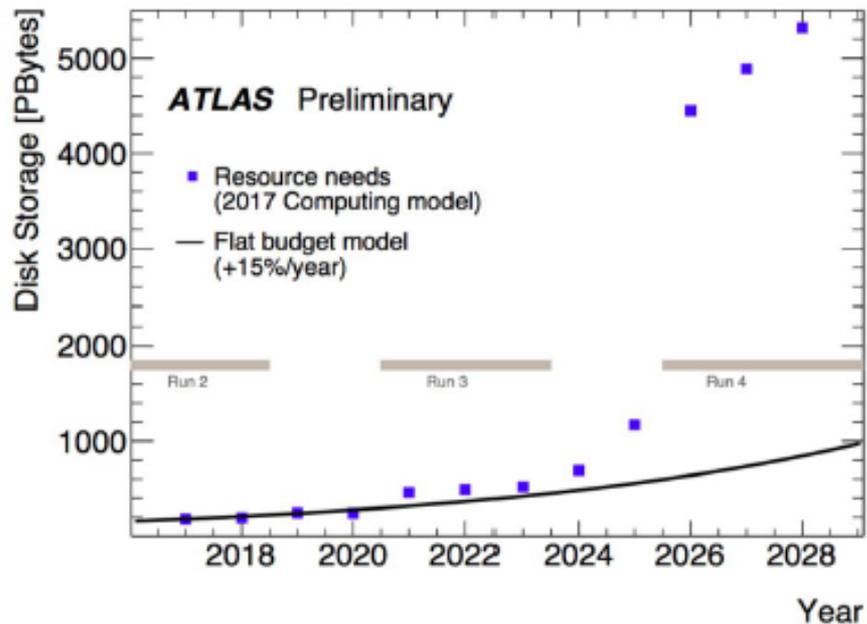
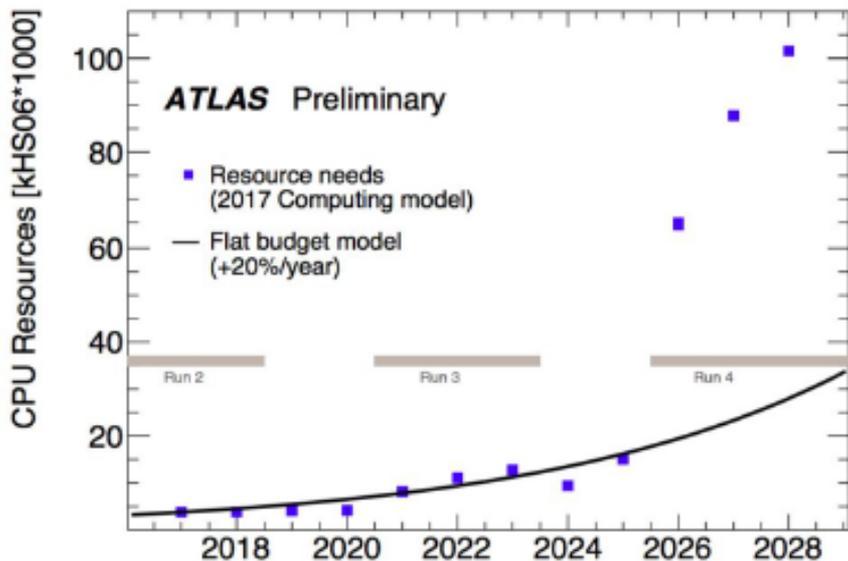
NSF DIBBs award #1724821, led by Rob Gardner UChicago

Equip the SciDMZ with a service orchestration platform to enable large-scale, multi-campus scientific compute services

Lessons Learned

- Dev & Ops cannot be separate silos
 - Early structure of ATLAS did not work well after LHC start
 - Slow dev projects were obsolete before deployment
 - Dev projects handed over to ops died quickly
- DevOps is not Dev+Ops
 - Middle stages of ATLAS worked better, but still too slow
- Teams need to work together continuously
 - DevOps is not Dev guiding Ops – leads to poor functionality
 - DevOps is also not Ops guiding Dev – software lacks design
- DevOps is participation and teamwork in endless loop
 - Where people are more important than tools and products

LHC Challenges Ahead



Conclusion

- DevOps is important for large scale data sciences
- Collaboration fosters good DevOps
- Culture is important – from the beginning

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development
(NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674,
Fax: 202-459-9673, Email: nco@nitrd.gov, Website: <https://www.nitrd.gov>

