



# BigPanDA for ATLAS

**Kaushik De**

**Univ. of Texas at Arlington**

**BigPanDA Workshop, ORNL**

**March 31, 2017**

# Overview




- **ATLAS has provided a testbed for BigPanDA development and commissioning for the past 4+ years**
  - First to run production in DOE HPC – Titan
  - Scaling up operations on Titan steadily over past 2 years
  - First to integrate network stack with workflow management system
  - Testing of BigPanDA packaging and distribution
  - Many other accomplishments – see Alexei’s talk yesterday
- **This is a natural role for ATLAS since PanDA grew up here**
- **ATLAS is unique in it’s ability to federate resources**
  - Started with grid resources – WLCG
  - Integrated cloud resources 5 years ago – research, Google, AWS...
  - Integrated HPC resources about 3 years ago – DOE, NSF, EU...
  - All of the federation is done through PanDA/BigPanDA

# Many Speakers Showed Similar Slide Yesterday




## Titan Core Hours Used by ATLAS



74 Million Titan Core Hours used in calendar year 2016  
In backfill mode – no allocation  
374 Million Hours remained unused.

# And Details of Backfill Like This







# The View from ATLAS (US Centric)



CPU consumption Good Jobs in seconds

52 Weeks from Week 00 of 2016 to Week 52 of 2016




Maximum: 17,168,416,337 , Minimum: 0.00 , Average: 9,267,960,602 , Current: 4,390,795,980

Full Geant4 Simulation  
Total CPU Usage by  
ATLAS in US for 2016




CPU consumption Good Jobs in seconds (Sum: 500,469,872,515)



Titan in Dark Green

Titan only runs G4 – other workflows not shown



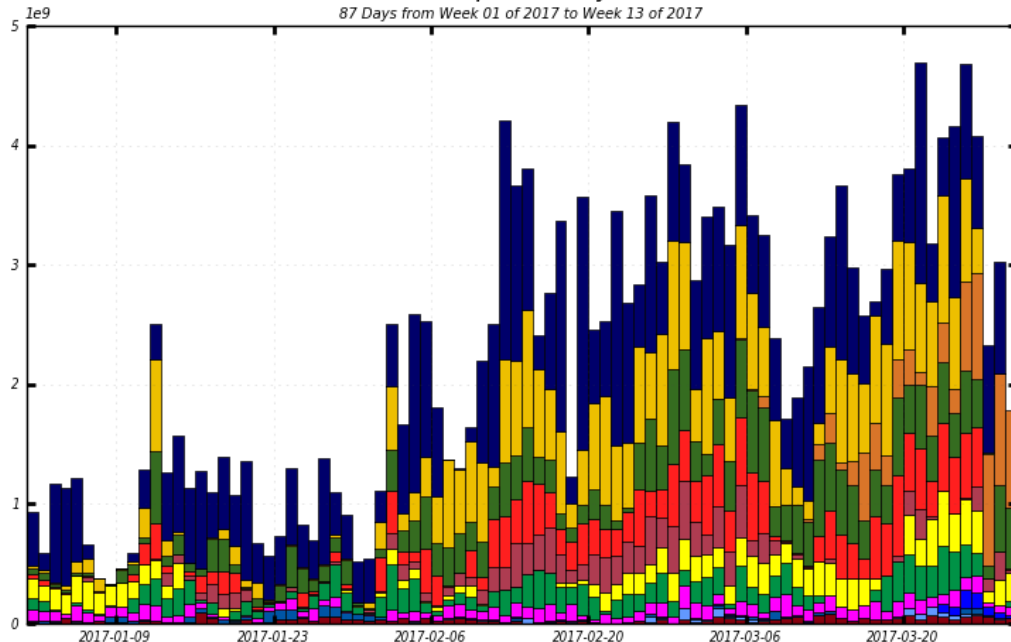
# And in 2017



dashboard

Wall Clock consumption Good Jobs in seconds

87 Days from Week 01 of 2017 to Week 13 of 2017




- OLCF
- WT2
- BNL-ATLAS-OPP
- BNL-ATLAS
- SWT2\_CPB
- HU\_ATLAS\_TIER2
- LBNL\_DSD\_ITB
- BU\_ATLAS\_TIER2
- LUCILLE
- MWT2
- UTA\_SWT2
- OU\_OCHEP\_SWT2
- AGLT2
- OU\_OSCER\_ATLAS
- WISC-ATLAS

Maximum: 4,696,155,743 , Minimum: 336,604,048 , Average: 2,250,169,472 , Current: 1,789,067,052

## Full Geant4 Simulation Total CPU Usage by ATLAS in US for 2017

shboard

CPU consumption Good Jobs in seconds (Sum: 500,469,872,515)  
OLCF - 15.95%



## Titan in Dark Green

Titan only runs G4 – other workflows not shown


- MWT2 - 18.96% (94,892,484,906)
- BNL-ATLAS - 15.42% (77,195,513,250)
- SWT2\_CPB - 10.43% (52,192,570,054)
- BU\_ATLAS\_TIER2 - 6.98% (34,924,397,648)
- BNL-ATLAS-OPP - 3.87% (19,384,574,100)
- HU\_ATLAS\_TIER2 - 1.69% (8,433,312,031)
- OU\_OSCER\_ATLAS - 0.12% (614,232,851)
- OLCF - 15.95% (79,816,098,098)
- AGLT2 - 10.46% (52,353,390,501)
- WT2 - 8.92% (44,642,019,931)
- UTA\_SWT2 - 4.24% (21,217,242,615)
- LUCILLE - 2.60% (12,999,128,216)
- OU\_OCHEP\_SWT2 - 0.36% (1,804,898,456)
- WISC-ATLAS - 0.00% (9,858)

# Working on Error Rates




2016

WallClock consumption in seconds



WallClock consumption in seconds




2017

# View from the Top of ATLAS



CPU consumption Good Jobs in seconds  
12 Weeks from Week 01 of 2017 to Week 13 of 2017



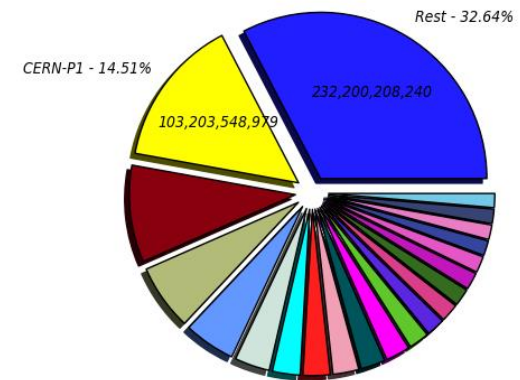
- Rest
- RAL-LCG2
- DESY-HH
- RAL-LCG2-ECHO
- GRIF-IRFU
- CERN-P1
- CERN-PROD
- FZK-LCG2
- SWT2\_CPB
- BOINC
- AGLT2
- LRZ-LMU
- TRIUMF-LCG2
- BNL-ATLAS
- MWT2
- BU\_ATLAS\_TIER2
- IAAS
- OLCF
- IN2P3-CC
- WT2
- TOKYO-LCG2

Maximum: 116,316,888,119 , Minimum: 0.00 , Average: 50,813,179,656 , Current: 45,438,683,489

## Full Geant4 Simulation Total CPU Usage by ATLAS for 2017



CPU consumption Good Jobs in seconds (Sum: 711,384,515,193)

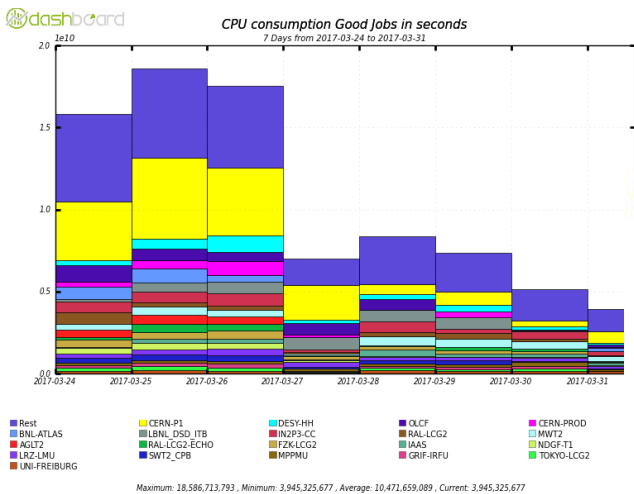


Includes all sites – not only US  
Titan only runs G4 – other workflows not shown


- Rest - 32.64% (232,200,208,240)
- BOINC - 9.44% (67,119,897,962)
- BNL-ATLAS - 5.00% (35,543,538,269)
- DESY-HH - 2.70% (19,192,257,140)
- LRZ-LMU - 2.53% (17,976,110,761)
- CERN-PROD - 2.36% (16,776,259,048)
- IAAS - 1.78% (12,644,179,789)
- BU ATLAS TIER2 - 1.67% (11,869,640,165)
- SWT2\_CPB - 1.58% (11,272,006,621)
- RAL-LCG2-ECHO - 1.42% (10,093,339,391)
- CERN-P1 - 14.51% (103,203,548,979)
- OLCF - 6.40% (45,539,830,926)
- MWT2 - 3.41% (24,269,479,923)
- AGLT2 - 2.63% (18,687,415,674)
- IN2P3-CC - 2.37% (16,847,988,523)
- RAL-LCG2 - 1.98% (14,057,105,186)
- FZK-LCG2 - 1.75% (12,423,983,538)
- TOKYO-LCG2 - 1.65% (11,707,424,824)
- GRIF-IRFU - 1.45% (10,294,630,663)



# Very Good Last Week



CPU consumption Good Jobs in seconds (Sum: 83,773,272,718)



Maximum: 18,586,713,793 , Minimum: 3,945,325,677 , Average: 10,471,659,089 , Current: 3,945,325,677

Rest - 31.06% (26,019,011,751)	CERN-P1 - 20.55% (17,211,823,692)	OLCF - 4.56% (3,816,628,246)
IN2P3-CC - 4.50% (3,768,091,881)	LBNL_DSD_ITB - 4.39% (3,678,036,281)	MWT2 - 3.74% (3,131,793,626)
DESY-HH - 3.70% (3,100,624,801)	CERN-PROD - 2.70% (2,260,654,233)	RAL-LCG2 - 2.61% (2,189,393,721)
FZK-LCG2 - 2.60% (2,178,977,088)	BNL-ATLAS - 2.59% (2,173,244,694)	LRZ-LMU - 2.55% (2,134,175,747)
SWT2_CPB - 2.00% (1,678,288,617)	IAAS - 1.95% (1,634,461,025)	RAL-LCG2-ECHO - 1.83% (1,535,361,270)
AGLT2 - 1.76% (1,470,543,058)	NDGF-T1 - 1.46% (1,224,989,387)	MPPMU - 1.41% (1,178,988,443)
TOKYO-LCG2 - 1.39% (1,161,851,478)	UNI-FREIBURG - 1.38% (1,154,391,364)	GRIF-IRFU - 1.28% (1,071,942,315)

# What can be done Better?



- Exploiting 1-2% of unused cycles on Titan provided about 5% of all ATLAS computing – same as large US sites
  - Great for Titan – improved efficiency at almost no cost
  - Great for ATLAS – huge impact on physics results
- But Titan has a factor of 2-3 more unused cycles
  - ATLAS cannot use them yet – waiting for Event Service
  - New PanDA server at OLCF can be used by others
- Plan for the next 6-12 months
  - Continue ATLAS usage of Titan – LHC is taking data
  - Many of the improvements discussed yesterday (and later today)
    - Harvester on Titan
    - Yoda on Titan
    - AES (using OS) on Titan