



HPCs in China

Wenjing Wu

Computer Center, IHEP, China

2017-9-22

Current HPC Resources

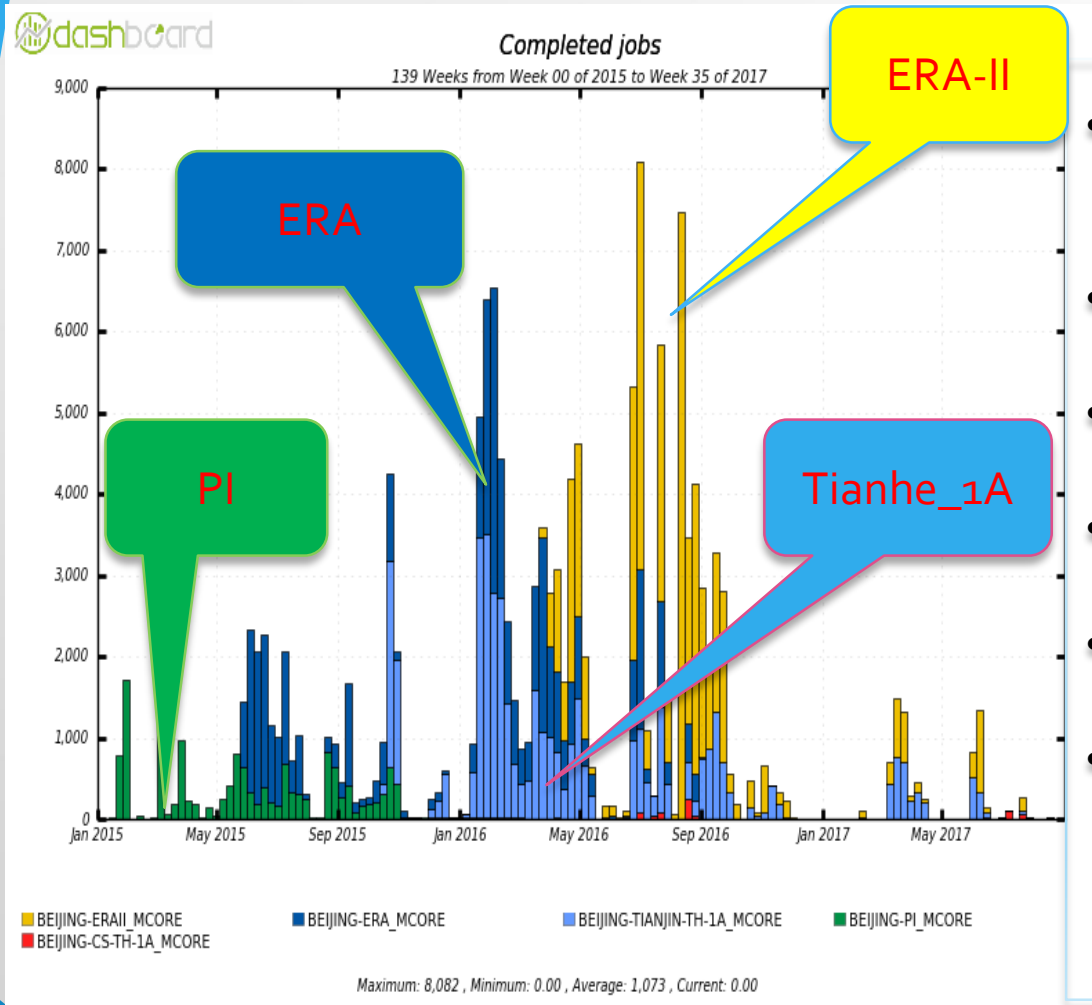
- 2 HPCs are available for use : Era-II and Tianhe_1A
- Provided through one Institute
 - CNIC(Chinese Network and Information Center, Chinese Academy of Sciences)
 - Common Interface to ARC CE: SCE (Science Grid Middleware)
- Current usable resource (Max. **2000** cores)
 - Era-II: Maximum 60 work nodes, **1440** cores
 - Tianhe_1A: Maximum 50 work nodes, **600** cores
 - Shared by different applications
- Can expect more (4000 cores?) with more effort

History

- Set up in Jan 2015 with PI (another HPC from SJU, China), started with one work node.
- March 2015, added Era (later upgraded to Era-II)
- Sep 2015, added Tianhe_1A
- Mar 2016, Era upgraded to Era-II
- Has some peak running between the middle of 2015 and 2016
- Became inactive due to little manpower support from ATLAS side to solve the technique issues since the end of 2016




Past usage




- Running Mcore jobs
 - 24 core (Era-II)
 - 12 core (Tianhe_1A)
- Used job cores grouped weekly
- 4 HPCs were used, 2 remain active
- Max. used of Era-II: **1400** cores/day
- Max. used of Tianhe_1A: **600** cores/day
- Currently active:
 - Era-II and Tianhe_1A
 - Both PI and Era stopped providing

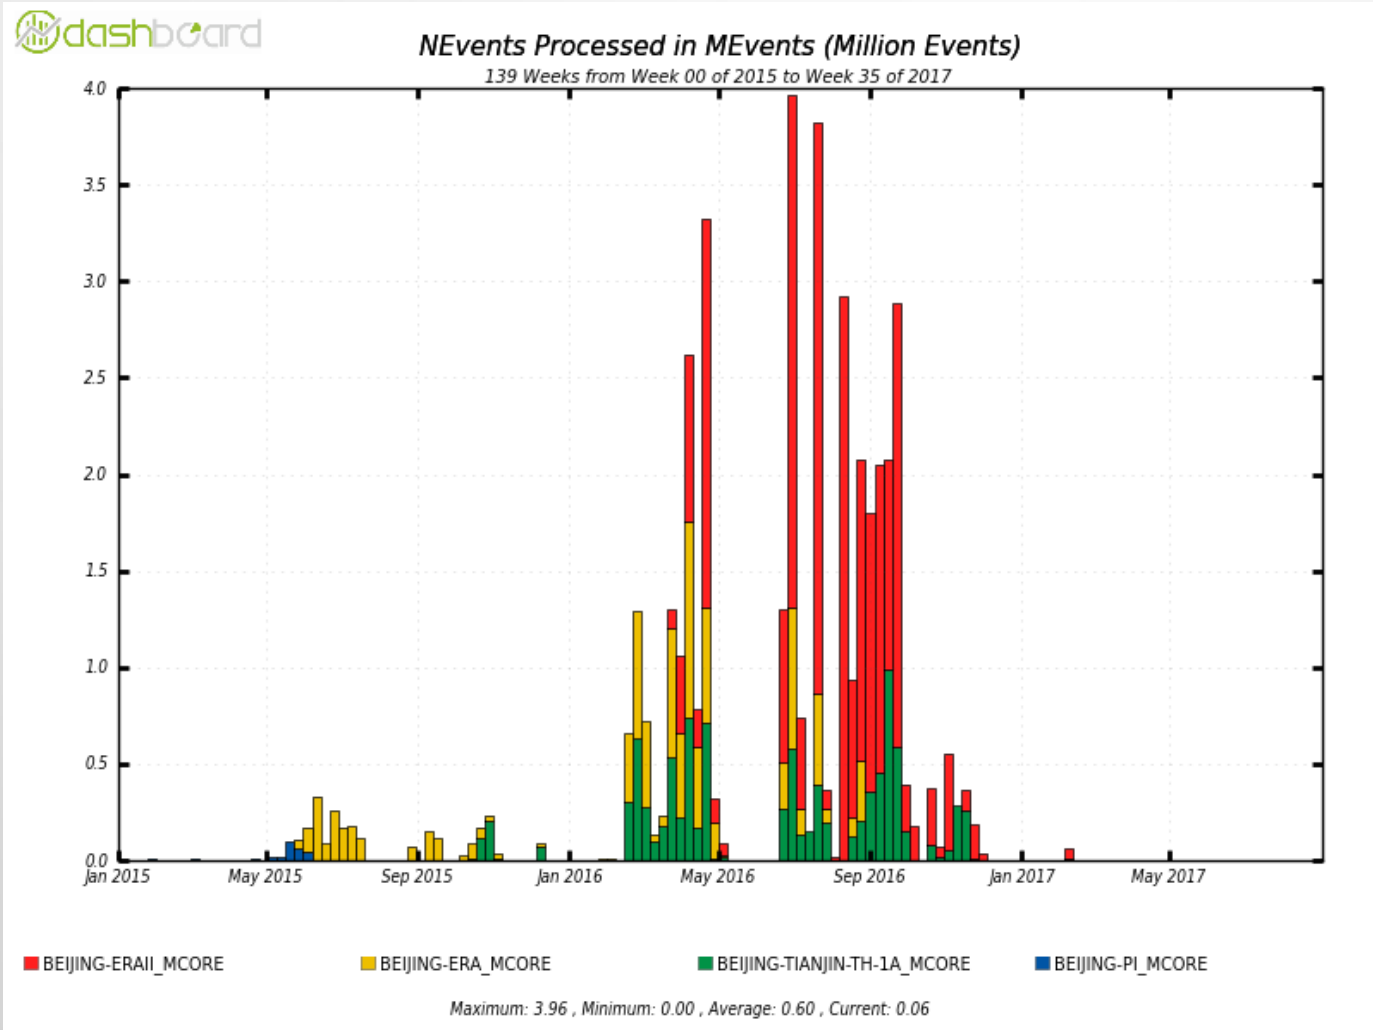
Peak used walltime (all jobs): **1700** CPU days/day



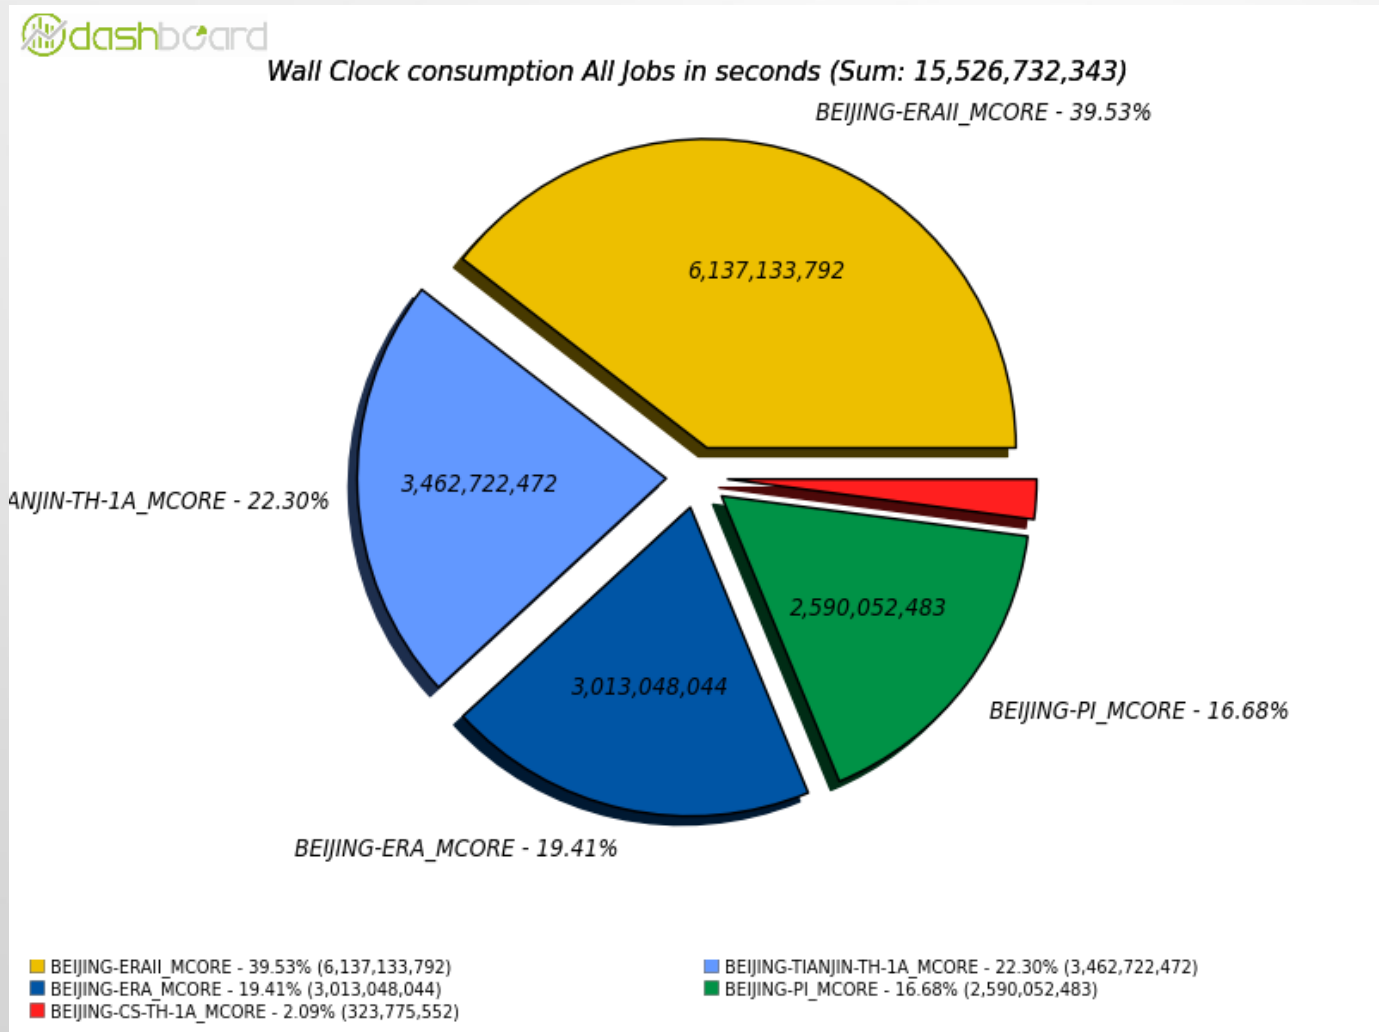
Peak used CPUtime (all jobs): **1500** CPU days/day



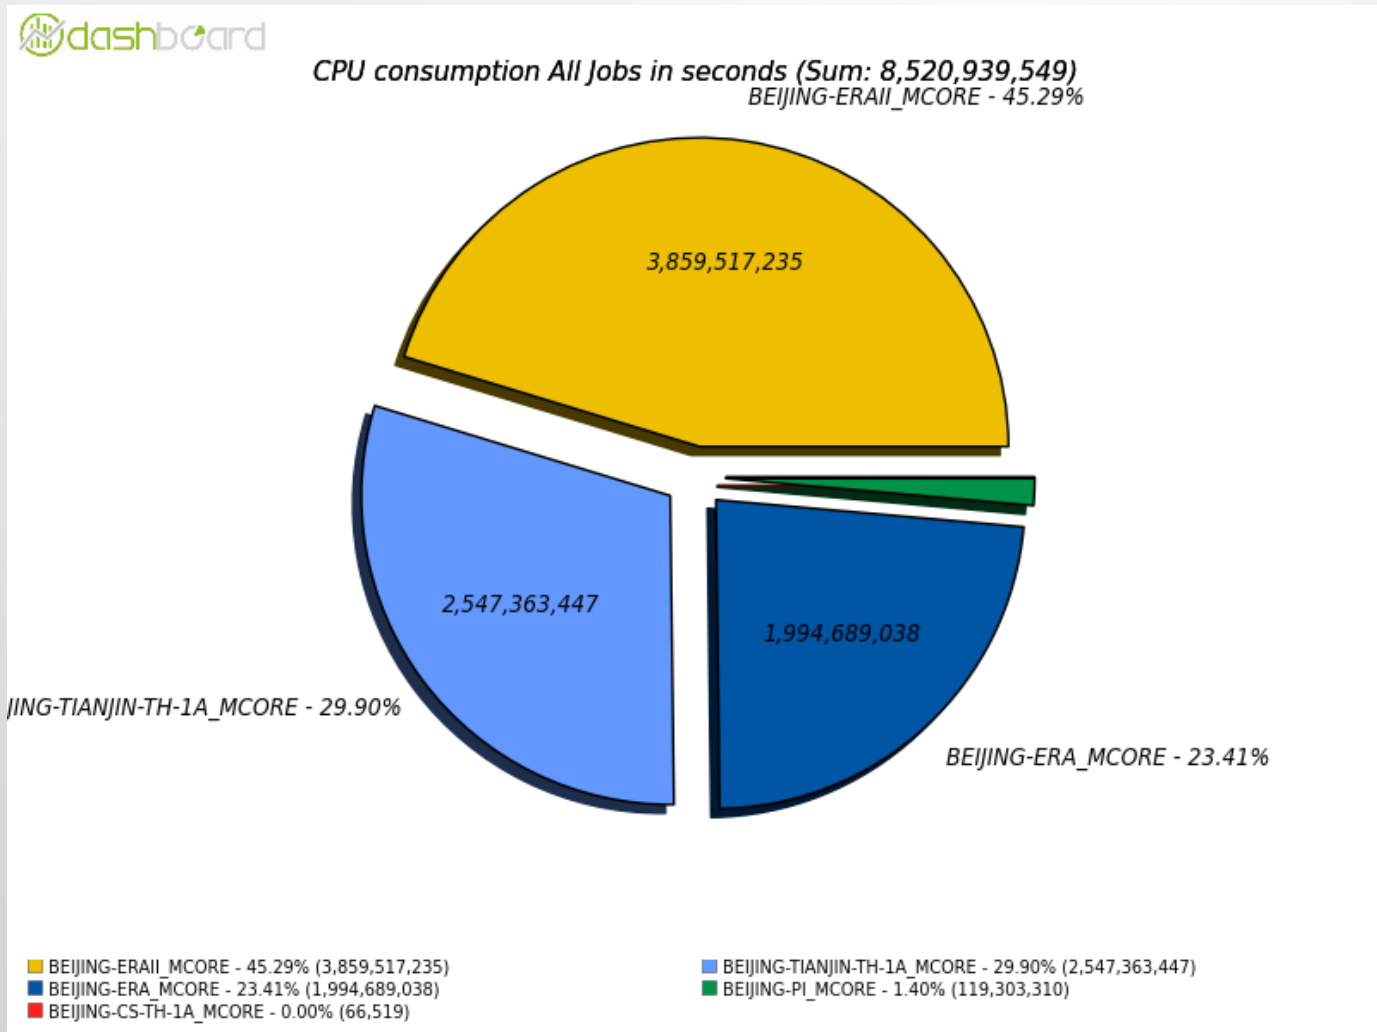
Peak simulated events (good jobs): **0.6M** events/day




All used walltime (all jobs): **492** CPU Years (since 2015)



All used CPUtime (all jobs): **270** CPU Years (since 2015)




All simulated events: **45 M events** (since 2015)



Errors




Number of Successful and Failed Jobs (Sum: 150,238)
Number of Failed Jobs - 47.24%



Panda Failures by ExitCode (Sum: 70,920)

General pilot error, consult batch log




Undocumented Execution Error Code: Exit Code: execution65

- General pilot error, consult batch log (43,682)
- Transformation not installed in CE (11,529)
- Undocumented Execution Error Code: Exit Code: execution68 (672.00)
- General exception raised which the transform did not handle (249.00)
- Undocumented Execution Error Code: Exit Code: execution38 (95.00)
- Get error: No such file or directory (52.00)
- Internal transform error (22.00)
- No reply to sent job (18.00)
- Failed during setup (13.00)
- Undocumented Execution Error Code: Exit Code: execution14 (8.00)
- Undocumented Execution Error Code: Exit Code: execution05 (13,098)
- Unspecified error: consult log file (745.00)
- lost heartbeat (442.00)
- DQ2 server error (115.00)
- Expired three days after submission (83.00)
- Transform received signal SIGTERM: _payload stdout file too big (32.00)
- Adder could not add files to the output datasets (21.00)
- Exception caught by pilot (10.00)
- Undocumented Execution Error Code: Exit Code: execution67 (8.00)

- 47% Failed jobs
- Due to
 - SCI interface to ARC CE, weak file transfer functionality
 - ATLAS Software release upgrade
 - Incompatibility from Pilot upgrade
 - Trying to access files from CVMFS
- Current obstacle
 - Distribute ATLAS R21.0 to HPC shared file system



Number of Successful and Failed Jobs
139 Weeks from Week 00 of 2015 to Week 35 of 2017



Can be Expected (1)

- Current Max. cores are **2000**
- All the resources are from the “Free” resources of HPCs, which is a small fraction of the entire resources
- These “Free” resources are shared by a few applications with equal priority.
 - Applications with more jobs can get a bigger share of resources
- Possible to add more from other Chinese HPCs
 - All the major Chinese HPCs are integrated into one interface (Science Grid), and managed by CNIC.
 - The technique challenges for extending to different HPCs are from share file systems, software distribution.
 - We did not progress well on this due to inactive follow-up and communication with the CNIC. (Manpower!)

Can be Expected (2)

- We can expect Max. ~ 4000 cores if
 - there is more manpower to solve the tech issues and remain good communication and connections with the HPC centers.
 - Motivate the HPCs centers to give more

Motivation for HPC centers

- They are not part of WLHC collaboration
- From HPC centers' point of view: "Free" resources are normally given to customers to "try", eventually they become paying customers. No HPCs are FREE!
- Political views: why should they give so much to ATLAS?
- They are interested
 - Visibility (paper, collaboration)
 - International collaborations which could bring them potential national/international funding opportunities
- Already done
 - CNIC people attended ATLAS software week in CERN in 2016
 - Attended CHEP2016, co-published paper in proceeding.

Complaints from HPC

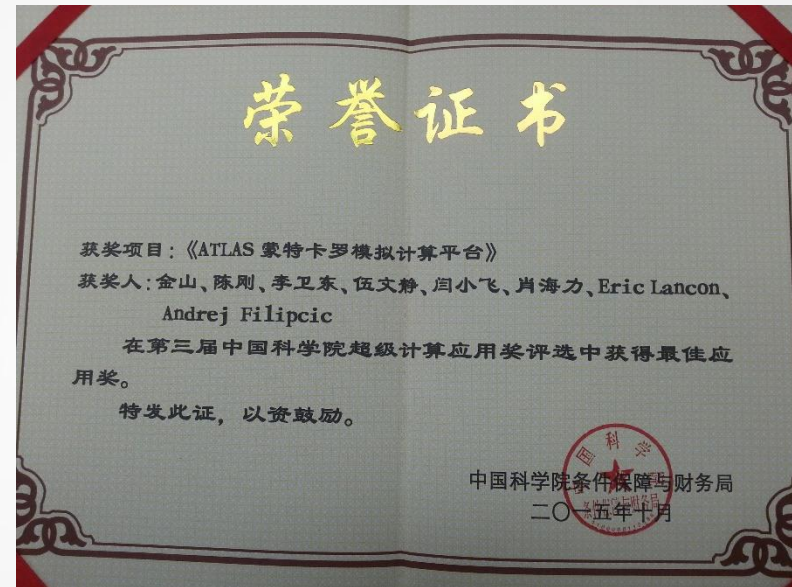
- Not fixing problems in time
 - Job keep failing for a few days, wasting CPU time
- Information is not transparent or being shared
 - Usage pattern (when and how much resources does ATLAS jobs need?)
 - Job monitoring information (how much are failing from ATLAS point of view)
 - The common ATLAS monitoring is a bit too big and complicated for the HPC center to look at, need to provided smaller/simpler job monitoring
- Solutions
 - Will add extra 0.25 FTE (Wenjing) in 2017 to work with HPC centers
 - monitoring running status, solving tech issues
 - bridging information sharing between HPCs and ATLAS
 - Motivation to HPC centers (common effort for funding search, publication)

Summary

- Chinese HPCs can provide a considerable amount of CPUs for ATLAS simulation.
- Was not fully used due to lack of manpower investment
 - Tech issues
 - Communication with HPC centers
- Hopefully can be improved in the later half of 2017
 - Reactive the usage of current available 2 HPCs (solve the new software distribution problem)
 - Add 1 or more HPCs to the Free resource pool

Acknowledgements

- ATLAS application won “Best Application” award in 2015 from CAS HPC alliance
- Ack. to
 - CNIC (HPC) : Xuebing Chi, Haili Xiao, Rongqiong Cao, Shasha Hou
 - NordGrid (Uni. Oslo) : s.h.raddum, Jon Kerr Nilsen, ATLAS : Eric Lancon, Andrej Filipcic, David Cameron, Paul Nilsson
 - IHEP : Gang Chen, Xiaofei Yan
 - SJU : Eric Germaneau





Thanks!