

OPERATION OF THE ATLAS DISTRIBUTED COMPUTING

F. Barreiro Magino¹, D. Cameron², A. di Girolamo³, I. Glushkov¹, A. Filipcic⁴, F. Legger⁵, T. Maeno⁶, R. Walker²

¹University of Texas at Arlington, ²University of Oslo, ³European Organisation for Nuclear Research (CERN), ⁴Josef Stefan Institute, ⁵Ludwig-Maximilians-University, ⁶Brookhaven National Laboratory

System Overview^[1]

- > 370 PB data (1.5 PB every week)
- > 180 sites in 40 countries
- GRID, HPCs, Clouds, Volunteers^[2]
- 12 GB/s global throughput
- 360k simultaneously running slots (up to 960k)
 - > 1.1M completed jobs per day (up to 2.7 M)




Fig.1 Amount of ATLAS data by 02.06.2018


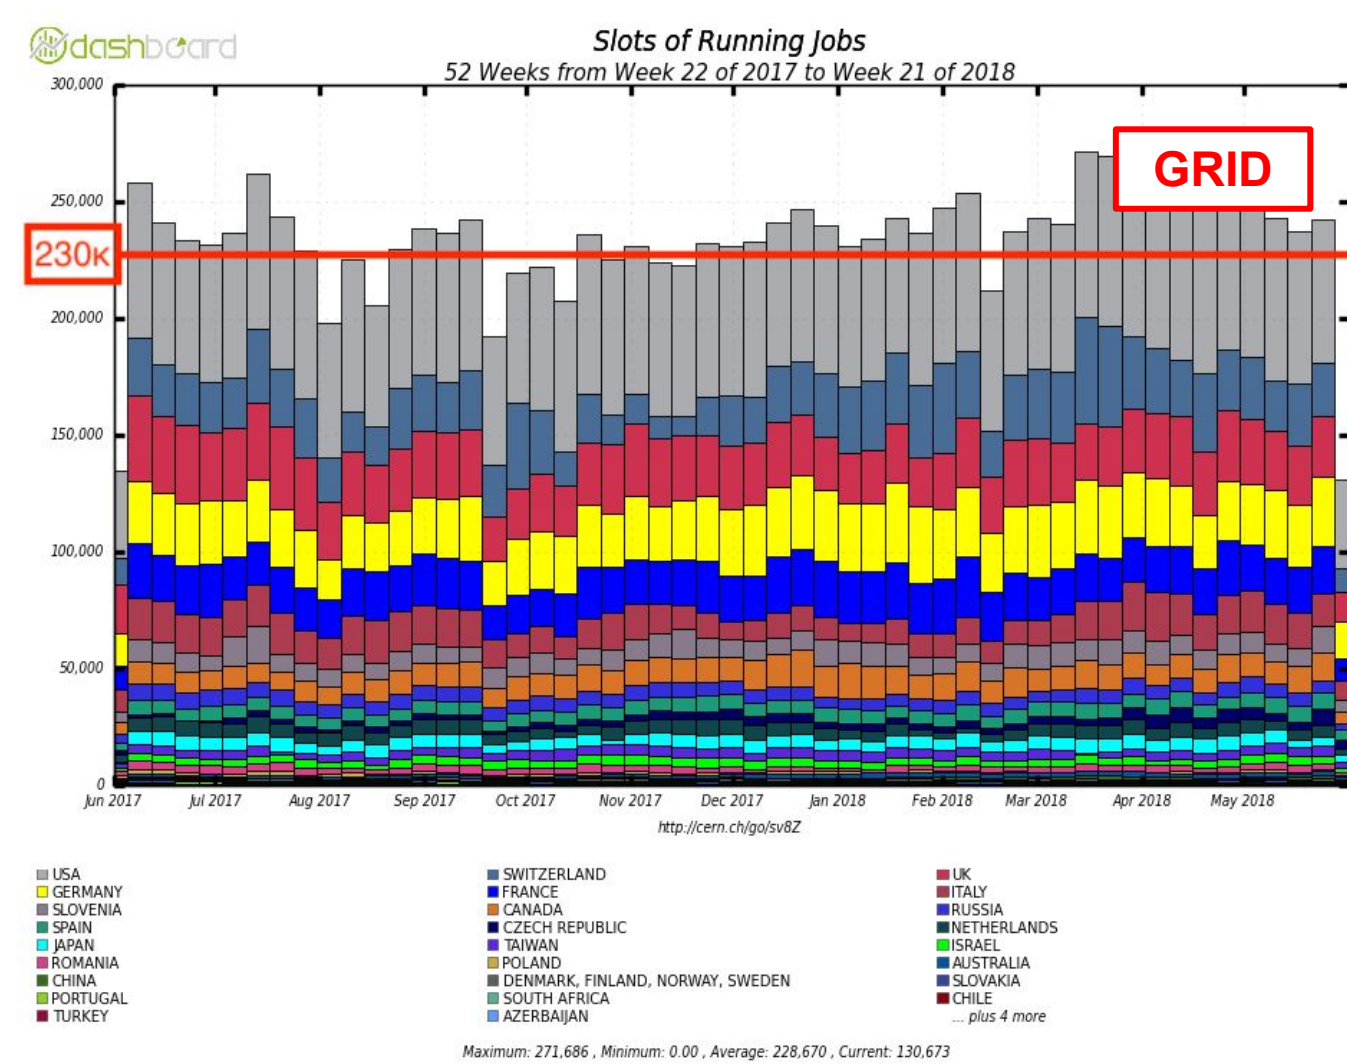


Fig.2 Running job slots contribution per resource type - GRID, HPCs, Clouds




Fig.3 Number of completed jobs per day grouped by Tier

Operations

- Manage expectations (workflow priorities) - Global shares^[3]
- Interaction with stakeholders (sites, analysis users, production managers)
- Accounting, monitoring and alarming
- System development and deployment

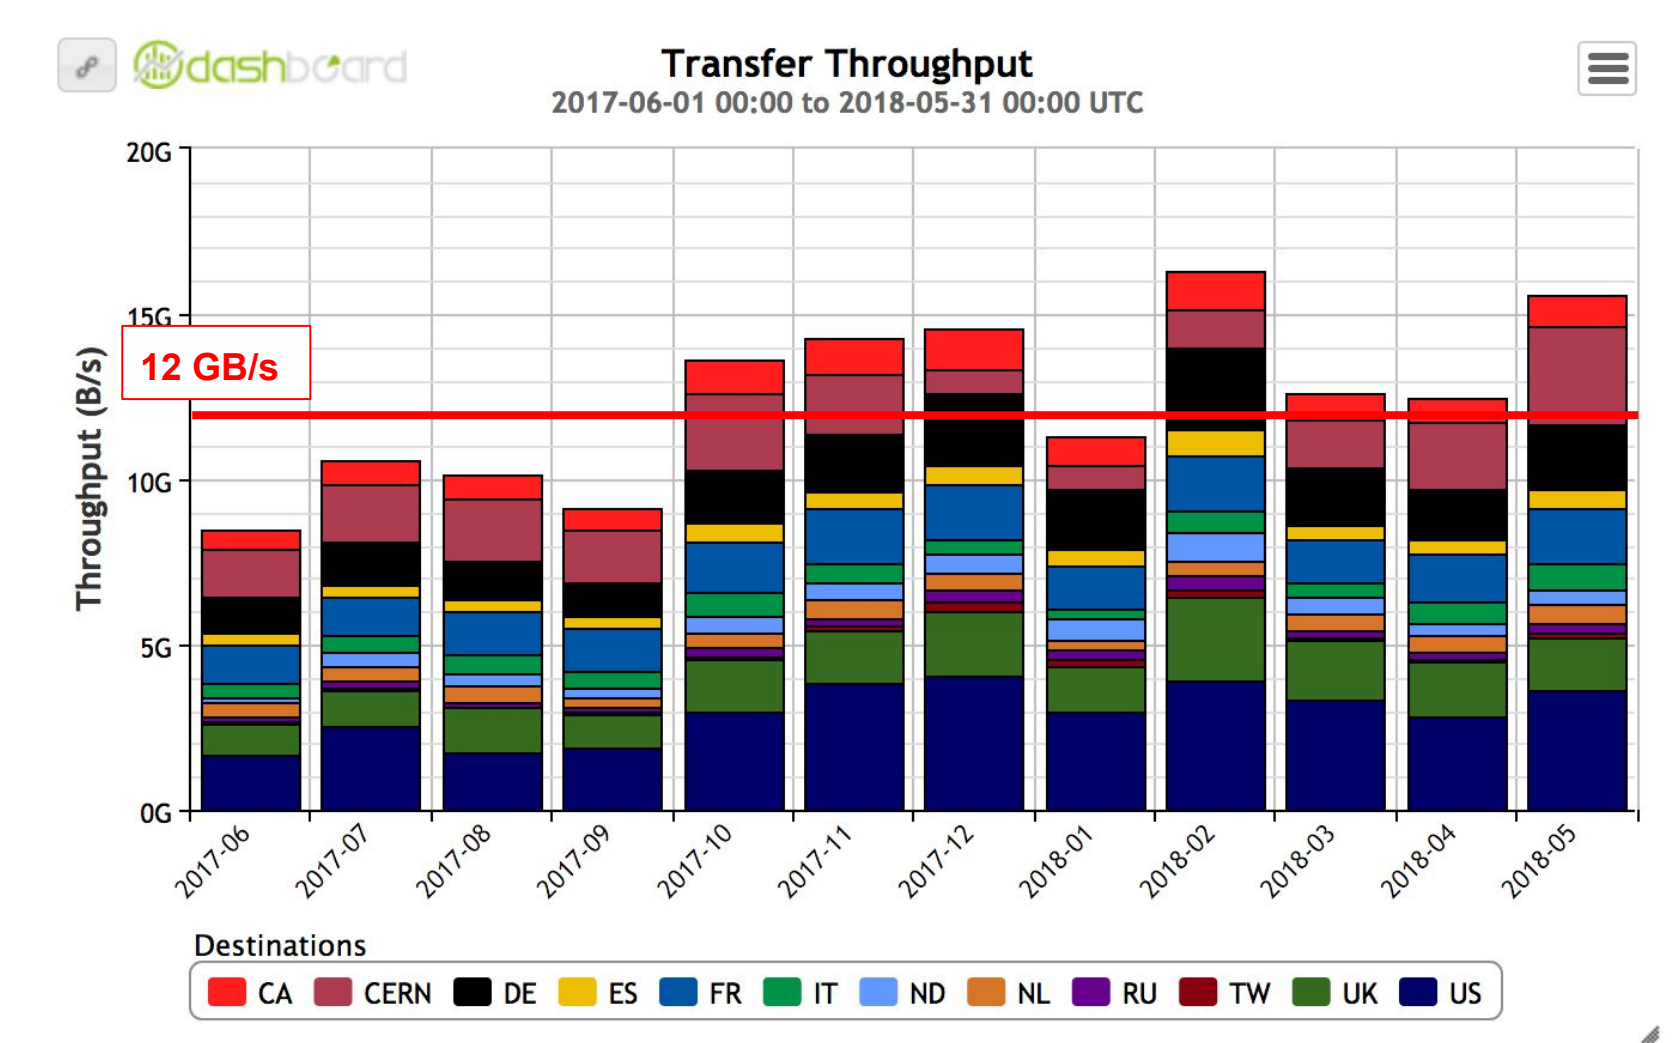


Fig.4 Average ATLAS data throughput in the last year

Optimizations driven by Operations

- Failure rate minimization
- Network utilization
- Disk and tape storage utilization
- Local batch resource utilization

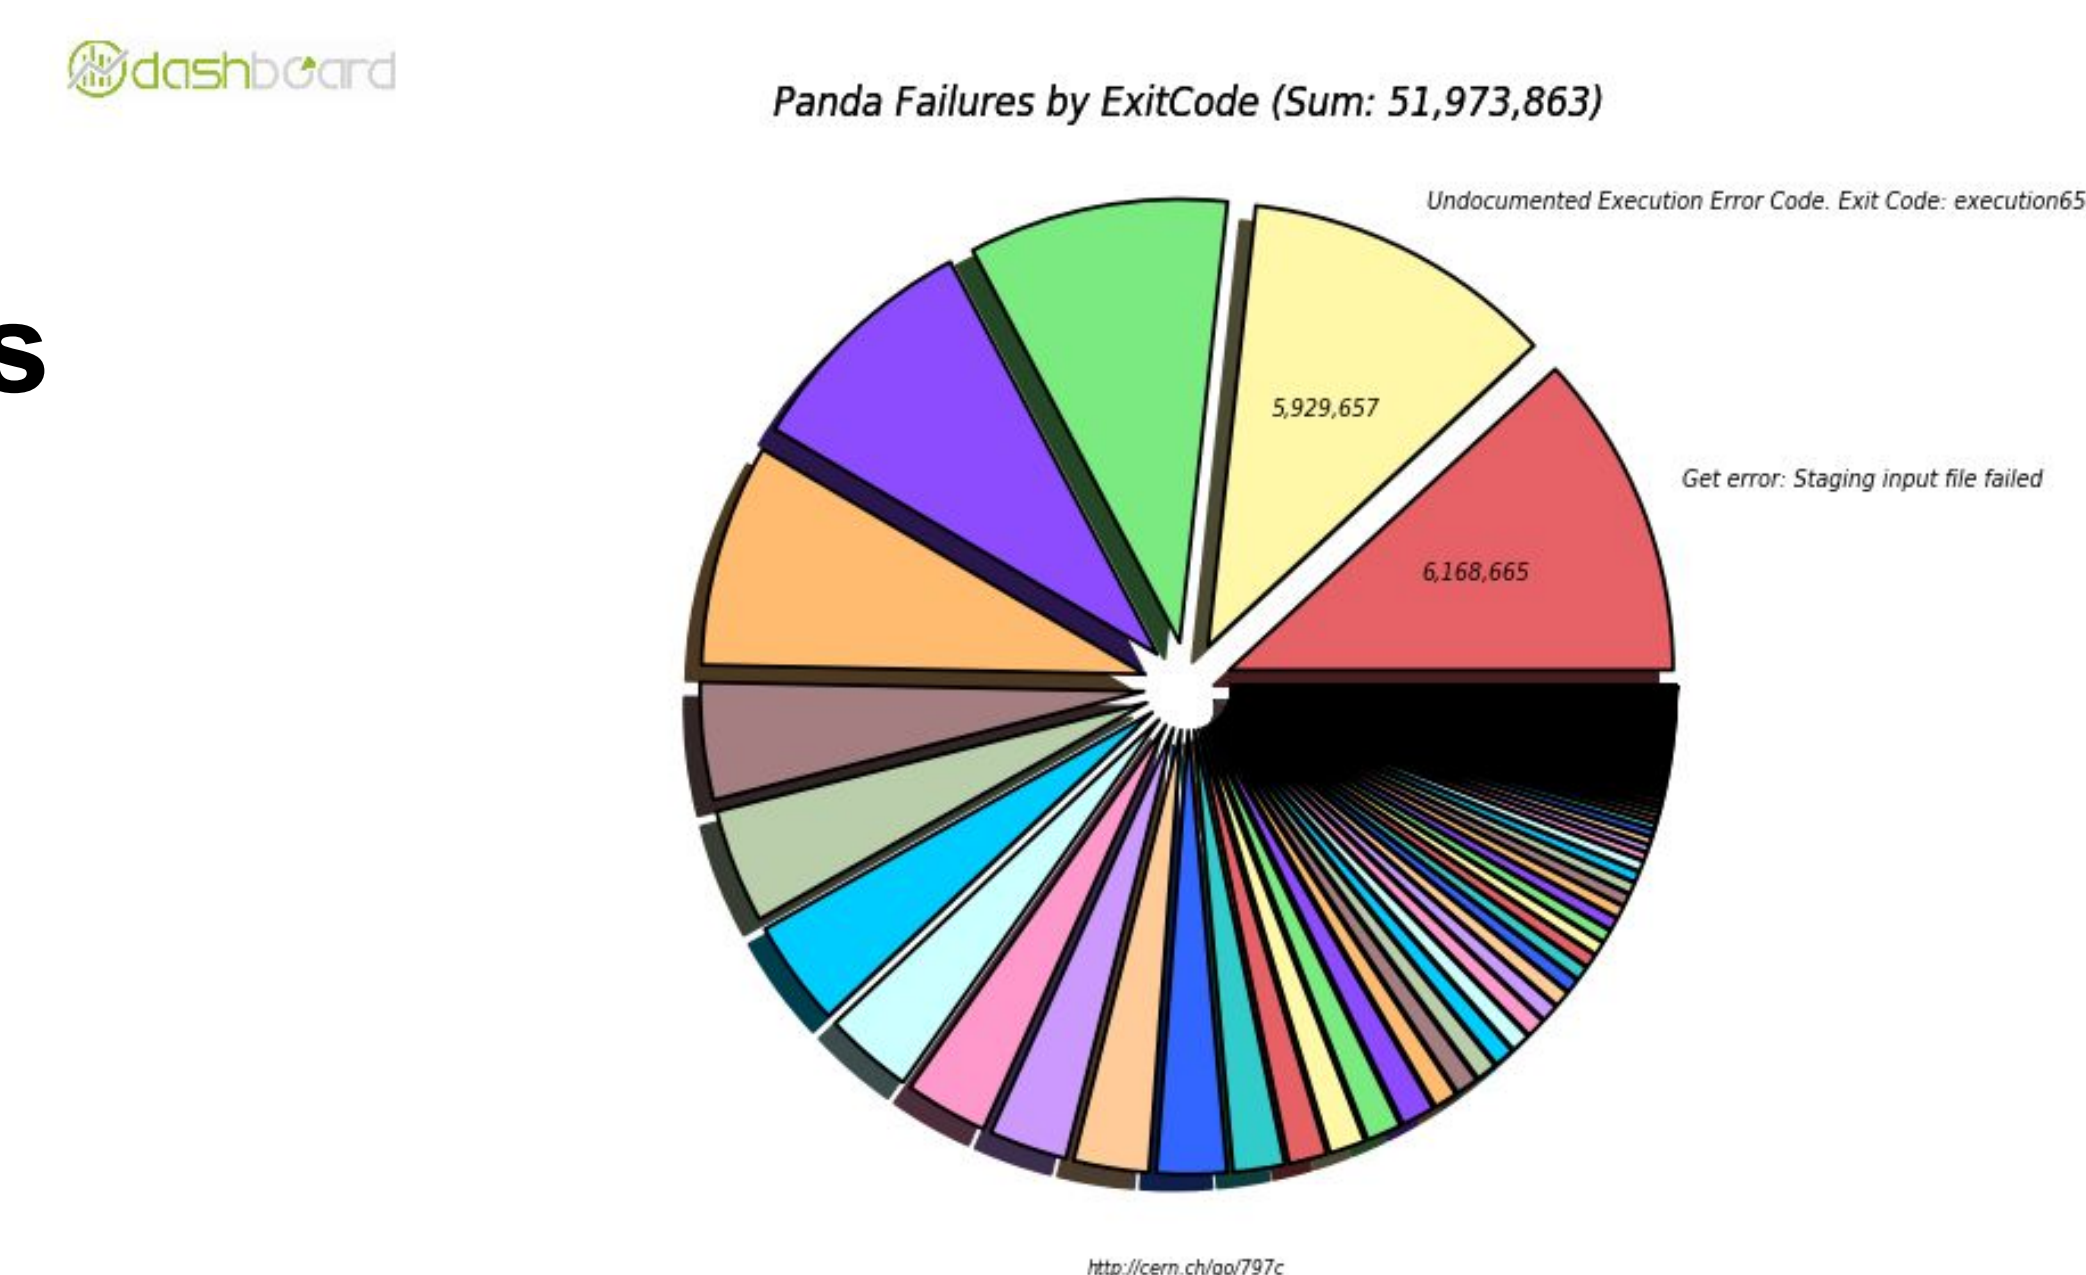


Fig.5 ATLAS Computing failures by category

Optimizations: I/O Intensity

- I/O intensity is defined as the input data size normalized to the job walltime
- It is heavily used in combination with the input data location and data size to be moved in the job brokerage.
- A central limit on both input data size to be moved and I/O intensity is preventing high I/O intensity jobs to be brokered on sites without the input data on the local storage and / or low network connectivity.
- Variations of the I/O limit is heavily affecting the global network utilization




Fig.6 Effect of I/O efficiency variations to the number of running job slots.




Fig.7 Effect of I/O efficiency variations on the production input transfer rate.

Optimizations: Panda Retrieval Module

- Each job can be retried n number of times, where n is defined by the user
- Retrieval module allows the operation team to limit or altogether forbid the retrieval based on the failure of a given retry.
- Web UI

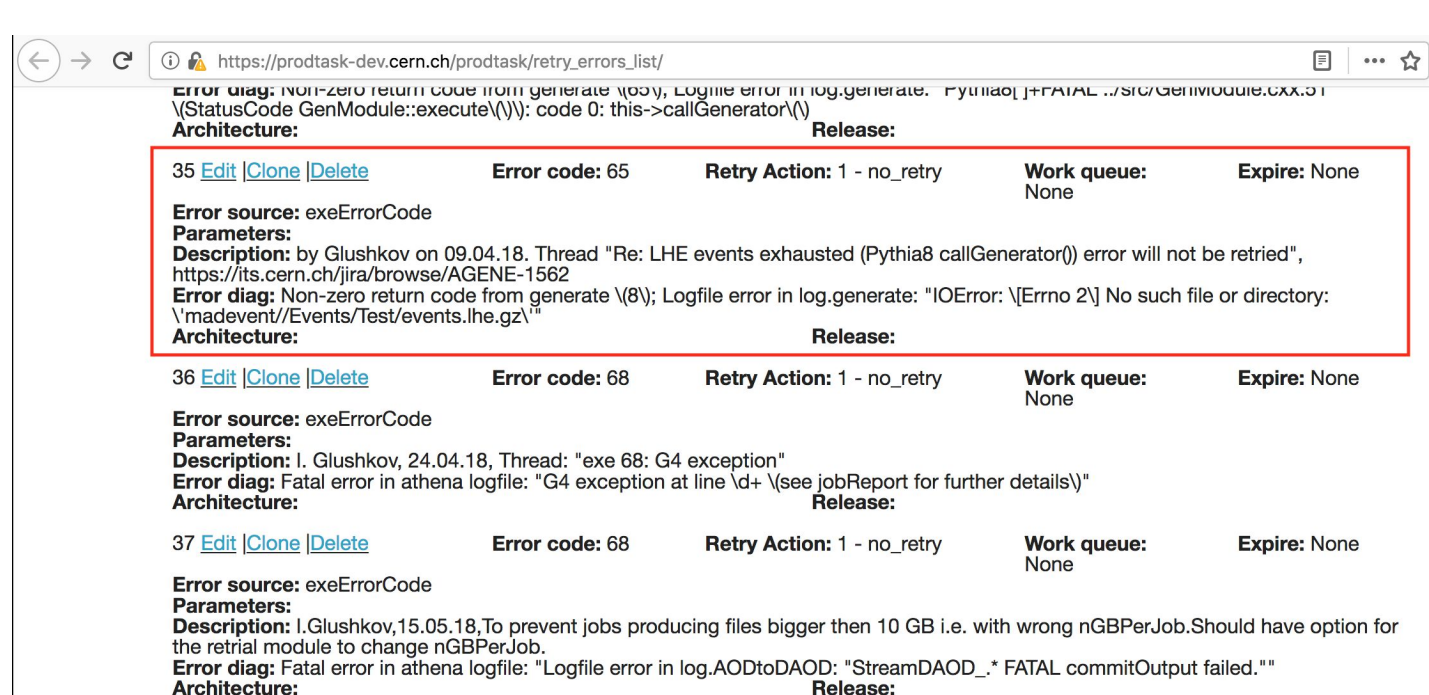


Fig.8 Web UI snapshot of Panda Retrieval Module Web UI




Fig.9 The impact of the introduction of a retrieval module rule over the corresponding failure's rate