



# ATLAS Grid Data Processing: system evolution and scalability

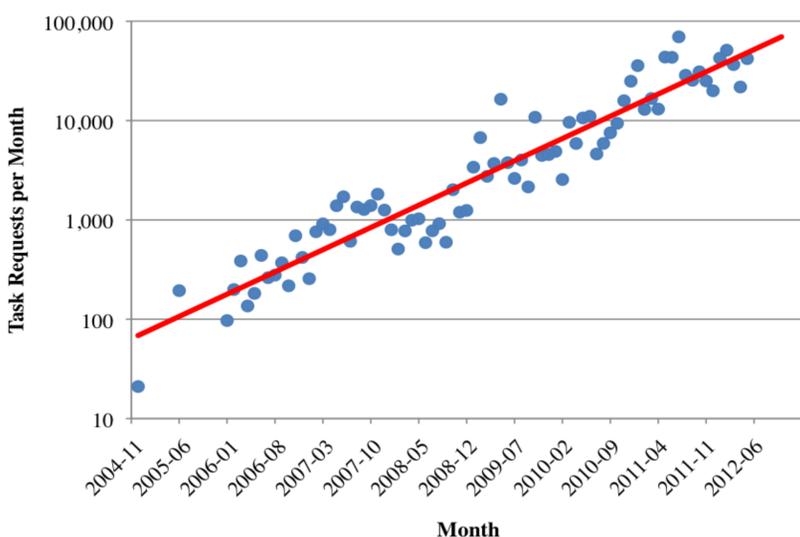


D Golubkov, B Kersevan, A Klimentov, A Minaenko, P Nevski, A Vaniachine and R Walker for the ATLAS Collaboration

**Abstract:** The production system for Grid Data Processing handles petascale ATLAS data reprocessing and Monte Carlo activities. The production system empowered further data processing steps on the Grid performed by dozens of ATLAS physics groups with coordinated access to computing resources worldwide, including additional resources sponsored by regional facilities. The system provides knowledge management of configuration parameters for massive data processing tasks, reproducibility of results, scalable database access, orchestrated workflow and performance monitoring, dynamic workload sharing, automated fault tolerance and petascale data integrity control.

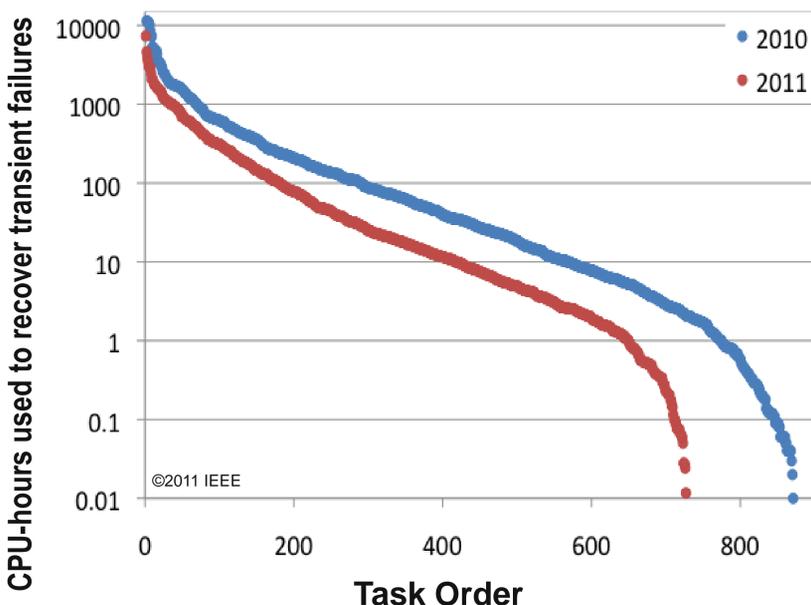
**Requirements:** The production system evolves to accommodate a growing number of users and new requirements from our contacts in ATLAS main areas: Trigger, Physics, Data Preparation and Software & Computing.

**Tasks:** In ATLAS production system the task (a collection of jobs) is used as a main unit of computation. The ATLAS job workload management system PanDA provides transparency of data and processing. Demonstrating system scalability, the rate of production tasks requests grows exponentially.



**Jobs:** Splitting of a large data processing task into jobs (small data processing tasks) is similar to the splitting of a large file into smaller TCP/IP packets during the FTP data transfer. Users do not care how many TCP/IP packets were dropped during the file transfer. Similarly, users do not care about transient job failures, when petascale data processing delivers six sigma quality performance in production, corresponding to event losses below the  $10^{-8}$  level.

**Performance:** Automatic job resubmission avoids event losses at the expense of CPU time used by the failed jobs. Distribution of tasks ordered by CPU time used to recover transient failures is not uniform: most of CPU time required for recovery was used in a small fraction of tasks. In 2010 reprocessing, the CPU time used to recover transient failures was 6% of the CPU time used for reconstruction. In 2011 reprocessing, the CPU time used to recover transient failures was reduced to 4% of the CPU time used for the reconstruction.



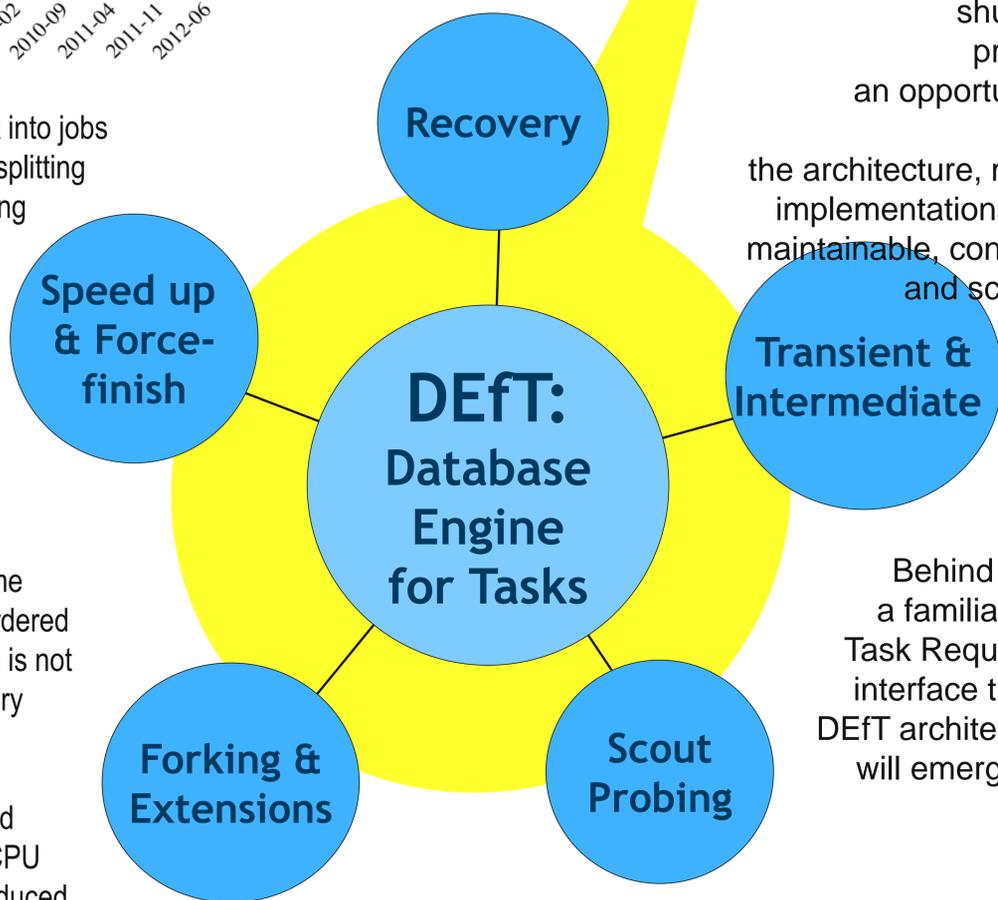
**Conclusions:** The GDP production system fully satisfies the requirements of ATLAS data reprocessing, simulations, and production by physics groups. The next LHC shutdown provides an opportunity for refactoring the production system, whilst retaining those core capabilities valued most by the users. We collected the requirements and identified core design principles for the next generation production system. Prototyping is in progress, to scale up the production system for a growing number of tasks processing data for physics analysis and other ATLAS main activities.

Tag Definition Task Request

Production system architecture grew organically. Legacy requirements constrained the design and implementation.

Task Evolution

The next LHC shutdown provides an opportunity to rethink the architecture, making implementations more maintainable, consistent and scalable.



Behind a familiar Task Request interface the DEFT architecture will emerge.

The design of the next generation production system architecture is driven by the following principles:

- ✓ *Flexibility:* As the requirements update process works well, the list of requirements and use cases continue to grow. Accommodating that growth, the next generation production system architecture must be flexible by design.
- ✓ *Isolation:* Avoiding inherent fragility of the monolithic systems, we adopted another core design principle - isolation. Separating core concerns, the production system logic layer will be separated from the presentation layer, so that users will have a familiar but improved interface for task requests.
- ✓ *Redundancy:* A core concern is scalability, required to assure eventual consistency of the distributed petascale data processing. Since in a scalable system, the top layers should not trust the layers below, we retain the redundancy.

