Automatic rebalancing of data in ATLAS distributed data management



Detecting imbalances





The ATLAS distributed data management system Rucio organizes over 220 petabytes of physics data across more than 130 sites worldwide. The aggregated transfer volume ranges between 1 and 2 petabytes every day.

ATLAS DDM

With the increased workload of Run 2 and the heavy utilization of data transfers, imbalances between computing slots and available capacity arise regularly. To compensate for this, the Balancer of Bits⁸ (BB8) daemon was developed.

- Semi-Automatic: In the semi-automatic mode the actual rebalancing is done automatically, but the event and volume have to be triggered by the operations team.
- Fully-Automatic: In the fully-automatic mode the daemon triggers rebalancing events for storage elements reaching a certain threshold automatically.



After the storage selection the workflow consists of two parts:

- **Data selection**: The algorithm selects datasets that are ATLAS official, have a long lifetime, are not popular (to minimize impacts on concurrent production), and do not have to be kept on the storage due to other reasons (rules).
- **Destination selection**: Key to the destination selection is the preservation of the original

Conclusion & Future work



The rebalancing daemon is used by ATLAS since

distribution policy. The algorithm looks up the replication policy for the original data placement and selects a storage element, out of this set, based on weighted selection on available space. The destination selection algorithm also includes a diminishing returns policy, to minimize ping-pong effects of rebalancing data between two storage elements.

May 2016 and already rebalanced over 2.5 petabytes of data.

Future work includes the possibility for BB8 to break policy borders, thus instead of rebalancing to a Tier-1 storage selecting a well-performing Tier-2 storage element. Also strategies to predict imbalances very early will be made to prevent them, before they happen.

