

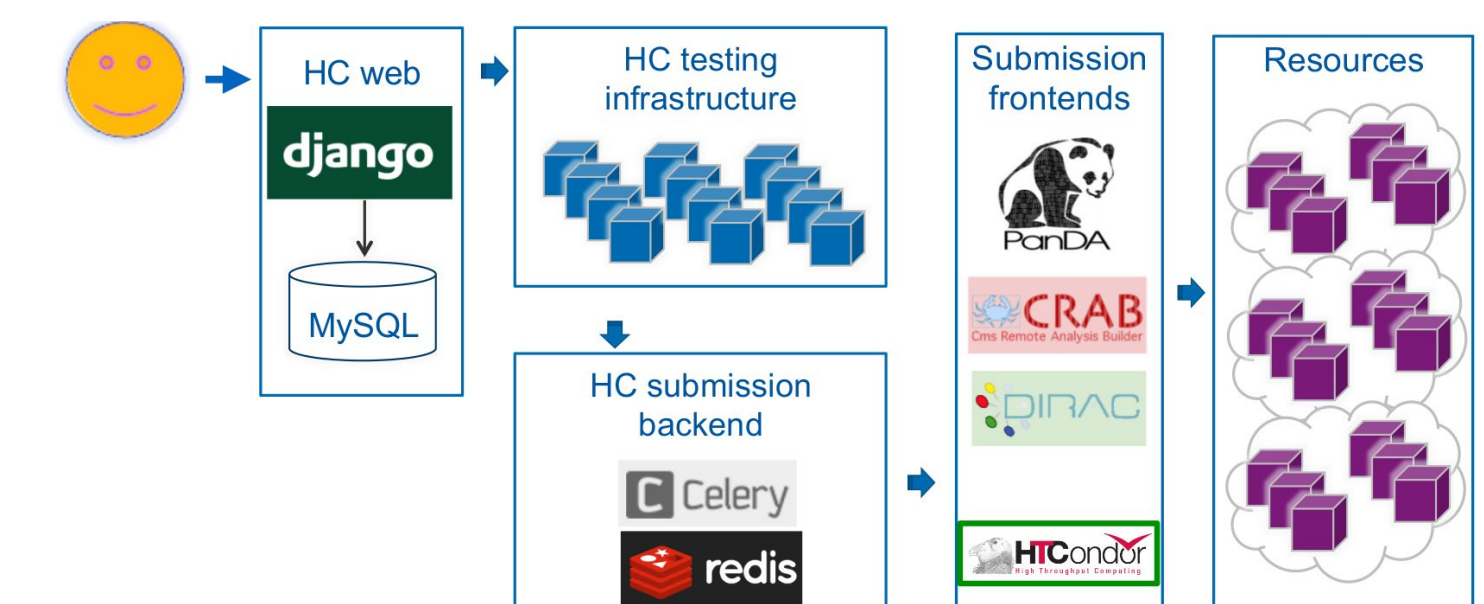
Improving ATLAS computing resource utilization with HammerCloud

J. Schovancová¹, F. Buehrer², J. Caballero Bejar³, G. Duckeck⁴, A. Fkiaras¹, F. Legger⁴,
T. Maier⁴, V. Mancinelli¹, G. Sciacca⁵, A. Yusta⁴ for the ATLAS Collaboration
1. CERN 2. Albert Ludwigs Universitaet Freiburg 3. Brookhaven National Laboratory 4. Ludwig Maximilians Universitat 5. Universitaet Bern

HammerCloud is a test framework.

... used to commission, test, and benchmark ATLAS Distributed Computing (ADC) resources and components of various distributed systems with realistic full-chain experiment workflows.

HammerCloud from far away



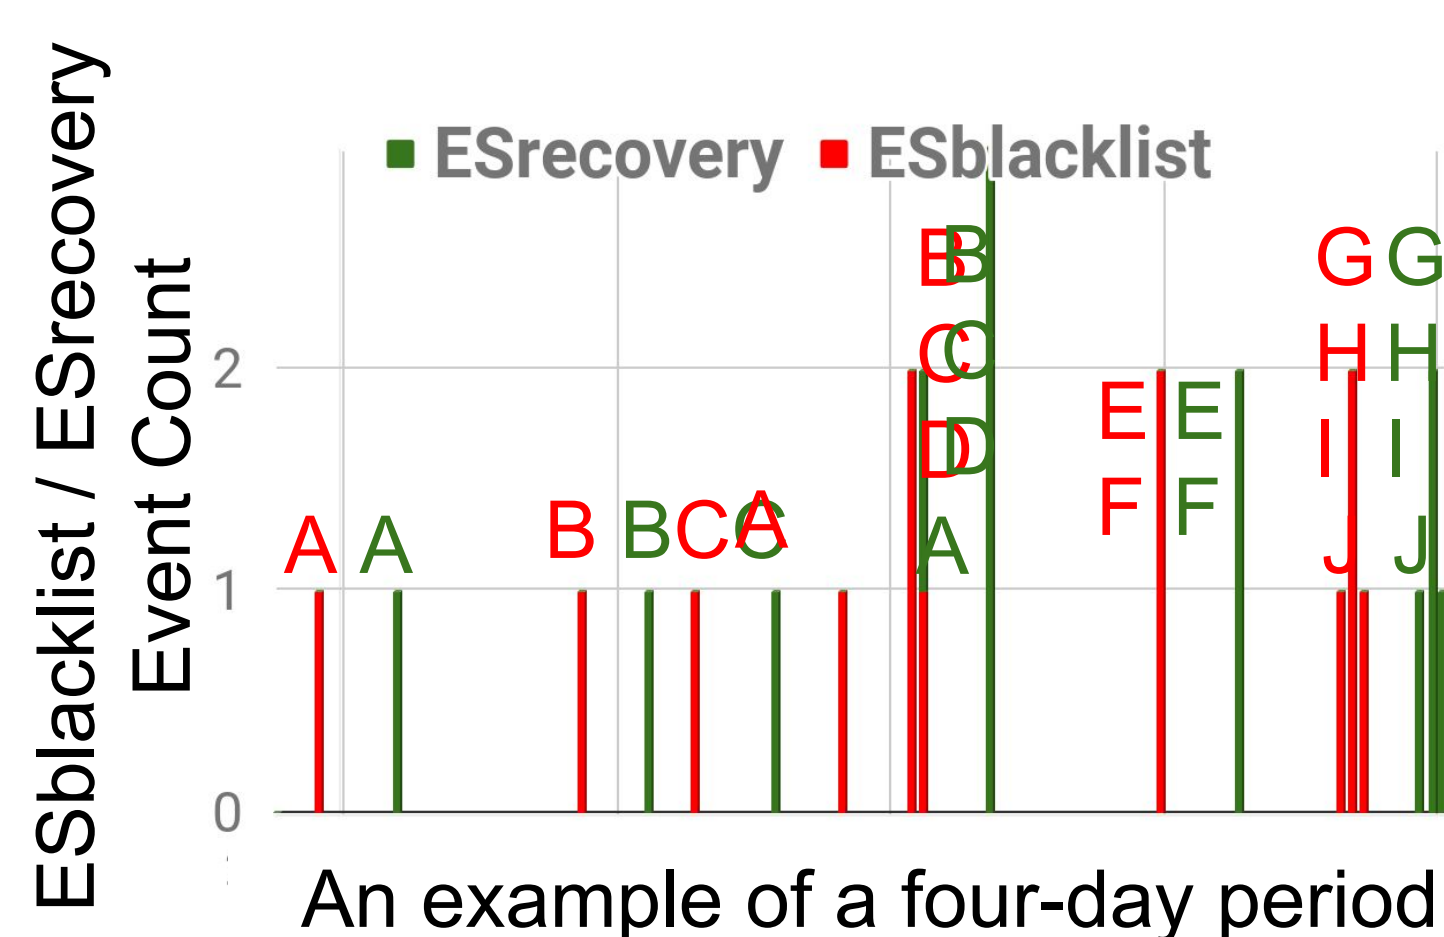
HammerCloud provides automated resource exclusion and recovery tools.

⇒ contributes to ADC Operations and Automation efforts.

HammerCloud evolved.

- Faster inclusion of new resources in tests. ... to maximize available compute cycles
- Auto-exclusion/-recovery of resources triggered by the ATLAS Event Service (ES)(*) job failures ... to help save manpower and speed up service recovery, and help ES experts to focus on other commissioning issues
- Prototyping machine learning algorithms for anomaly detection. ... to get rid of false positives in blacklisting
- Functional & stress testing of Object Stores ... to help with commissioning & operations of the Event Service infrastructure

(*) ATLAS Event Service has been deployed to exploit diverse resources that can disappear on short notice. It assigns event-level granularity workloads to running application, and streams outputs away almost continuously.



ESblacklist/ESrecovery events of different queues (A-J) in time.

ESblacklist: >50% jobs ESfailed in the past 2 hours
⇒ *disable queue for ES jobs*

ESrecovery: re-enable ES after 6 hours

EXHAUSTED: ESblacklist 4x within 24 hours
⇒ *no recovery, alert ES experts*

HammerCloud keeps helping with commissioning.

... besides commissioning of new resources, we help with commissioning of concepts and components of distributed systems:

- UnifiedQueues: Simplified configuration of queues for workflows of different activities
- Components of PanDA pilot: new data movers
- AGIS controller: component of ATLAS Grid Information System (AGIS) responsible for: providing an API to dynamically set status of queues by different entities and probes
- Components of Rucio: helped to clean up a data access protocol zoo, commission direct data access

