



Argonne
NATIONAL
LABORATORY

... for a brighter future



U.S. Department
of Energy

UChicago ►
Argonne_{LLC}



A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC


ATLAS Metrics for CCRC'08 Database Milestones

*WLCG CCRC'08 Post-Mortem Workshop
CERN, Geneva, Switzerland
June 12-13, 2008
Alexandre Vaniachine*




3D ATLAS Milestone Accomplishment

- 3D database services at Tier-1 sites are critical for data reprocessing
 - Since data reprocessing depends not only on database access it was launched at the end of the CCRC'08 exercise
- 3D servers at the Tier-1 sites were asked to operate at the milestone metrics, which is the sessions load from ATLAS data reprocessing jobs
- In a steady-state LHC data processing operation we require capacities capable of sustaining ~1k concurrent Oracle sessions in total
 - **During CCRC'08 3D services demonstrated capability to exceed that cumulative ATLAS Milestone target by 1/3**
- That milestone was achieved, with some of Tier-1 sites not yet tested:
 - Six sites operated above the ATLAS milestone metrics
 - *two sites operated close to $\frac{3}{4}$ of the metrics*
 - *two remaining sites were tested at less than $\frac{1}{2}$ of the metrics*
- **Based on the CCRC'08 experience to date, we are confident that critical WLCG 3D services are ready for ATLAS LHC data processing**





What Else Has Been Accomplished During CCRC'08

- 3D streamed to Tier-1 sites all the COOL data needed for bulk reprocessing
 - All ATLAS data reprocessing jobs got the data replicated to Tier-1s
- Recently upgraded RAC hardware at CERN delivered enough capacities to handle unexpected workload from FDR-2 data processing at Tier-0
- At the tested Tier-1 sites reading of the streamed DCS data using latest COOL 2.4 software release worked with good performance
 - The DCS data need special attention because of the largest volumes
- “Upstream” 3D streams replication has been setup from Muon Calibration Centers in Rome, Michigan and Munich to INTR server at CERN
 - Used during ATLAS FDR-2 exercise during CCRC'08
 - *Critical for ATLAS “24-h calibration loop”*
- Database monitoring worked well



Monitoring Burst Load at Tier-1s



TRIUMF
CPU load at peak:
■ 98% peak for 1 min
■ 70% average



Short burst periods correspond to submission of limited number of test jobs for data reprocessing tasks

Note: Summary load on both Oracle RAC nodes at TRIUMF Tier-1

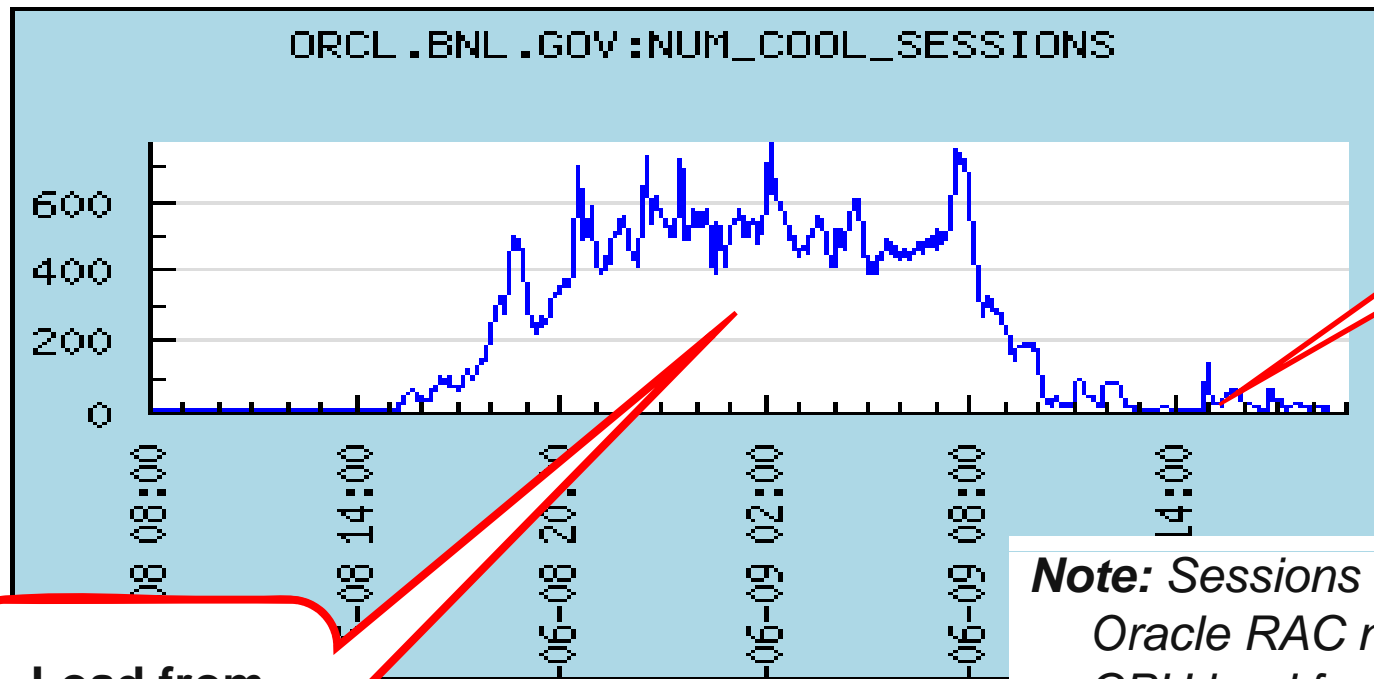
- ATLAS does not have spare WLCG 3D Oracle capacities deployed to support burst loads when many data reprocessing jobs start at once on an empty cluster



Monitoring CCRC'08 Bulk Data Reprocessing at Tier-1





Monitoring CCRC'08 Cosmics Data Reprocessing at Tier-1*



Load from US ATLAS users work


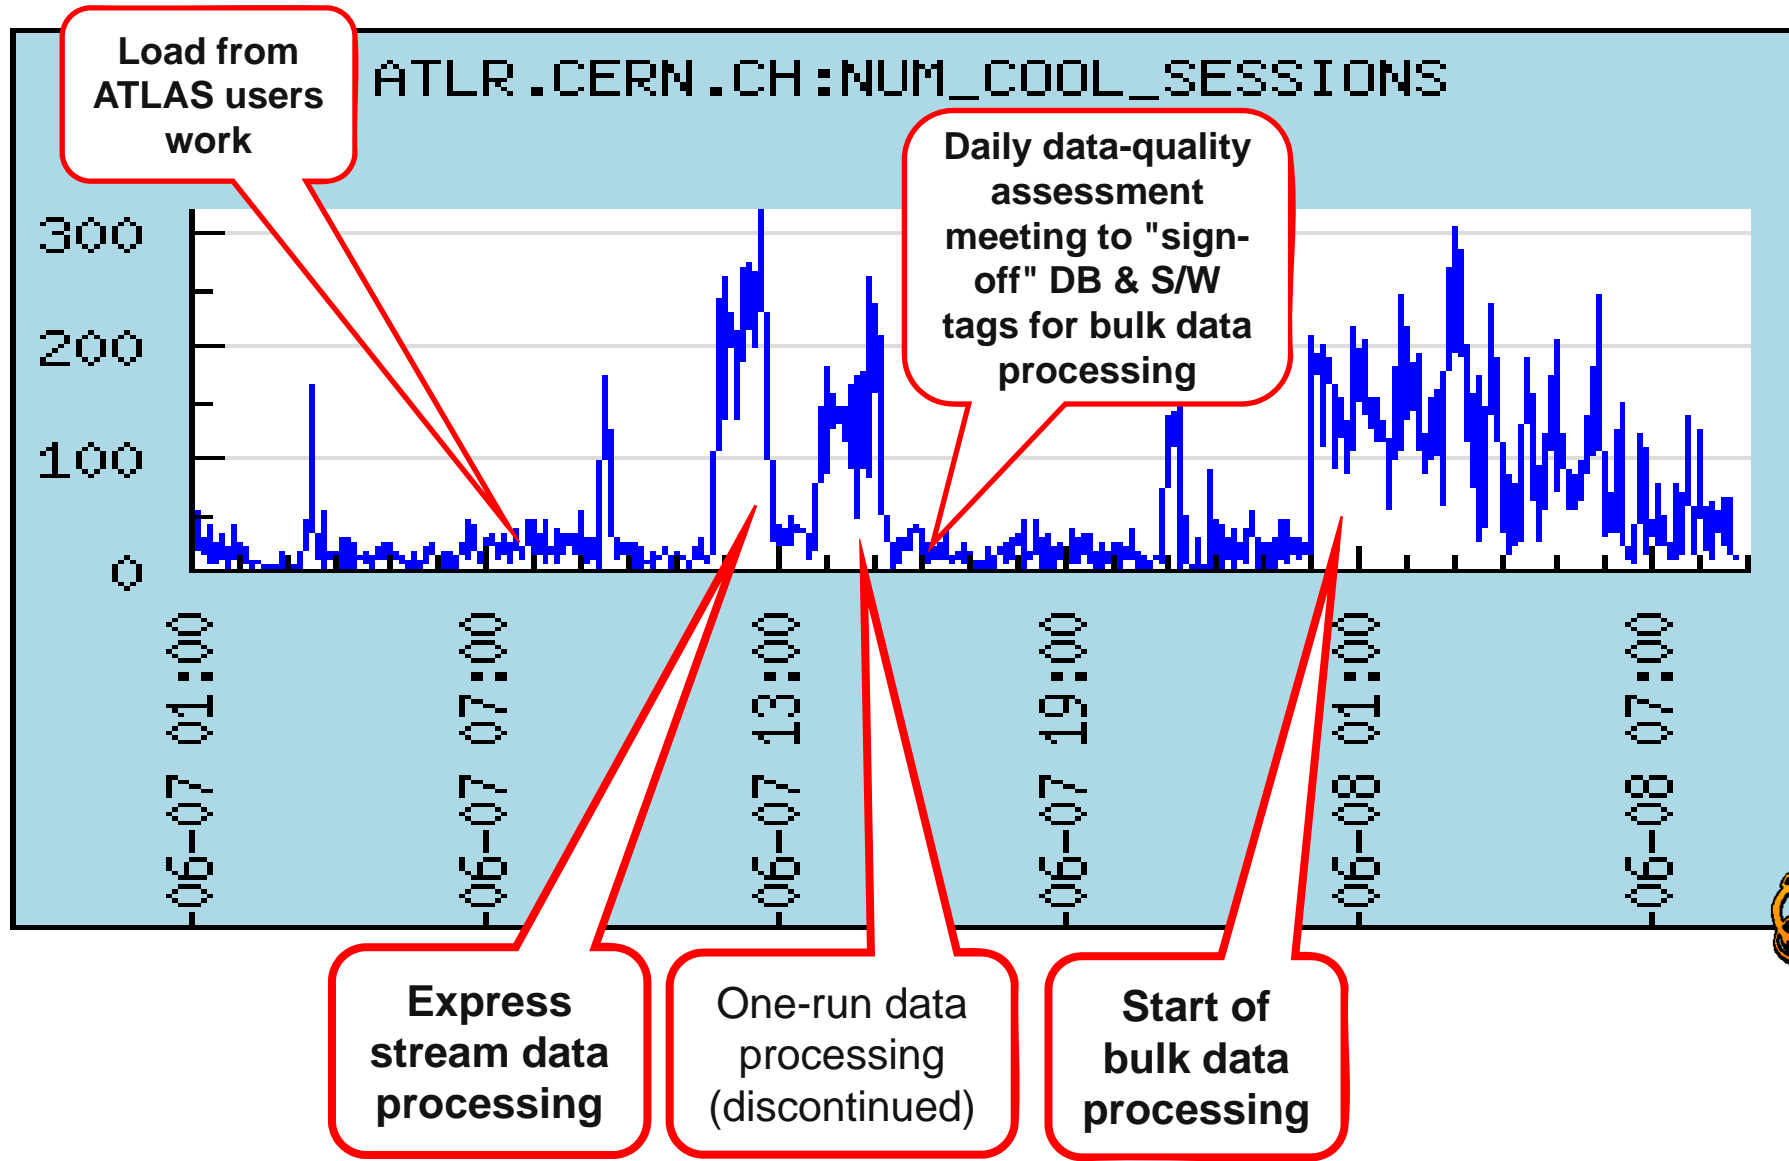
Load from bulk data re-processing

Note: Sessions load from both Oracle RAC nodes at BNL Tier-1, CPU load from a single node




* Reprocessing using large ATLAS Tier-2 sites in the U.S. in addition Tier-1 resources at BNL

CCRC'08 Monitoring of FDR-2 Data Processing at Tier-0



DCS Scalability Test Load at CNAF Tier-1



30% CPU load at max sessions shows that ATLAS capacities request was not excessive


ATLAS Conditions DB scalability tests at CNAF with new COOL 2.4 software show good results for large DCS data reads

Note: Summary load on three Oracle RAC nodes at CNAF Tier-1



CCRC'08 Post-Mortem Findings

- Latest scalability test results at CNAF show that WLCG 3D capacities deployed on ATLAS request are not excessive
 - For example, ATLAS does not have spare WLCG 3D Oracle capacities deployed to support burst loads when many data reprocessing jobs start at once on an empty cluster
- Oracle connect strings changes should be planned and announced well in advance, so that they can be propagated to ATLAS jobs configurations
- Evaluation of bulk data reprocessing at the Tier-2 sites uncovered new requirements:
 - Deployment of additional Oracle hardware capacities at BNL Tier-1 may be required to support bulk reprocessing of LHC data at large Tier-2 sites in the U.S.
 - How to connect to remote Oracle server via the site firewall?



Conclusions and Next Steps

- The WLCG 3D services deployed for ATLAS demonstrated capability to exceed cumulative ATLAS Milestone target by 1/3
 - Based on the CCRC'08 experience to date, we are confident that critical WLCG 3D services are ready for ATLAS LHC data processing

- Since there is little room for an error in our estimates of 3D capacities required for ATLAS
 - We have to validate these estimates using latest ATLAS software and Computing Model

- ATLAS validation of 3D services will continue with further data reprocessing exercises increasing in scope and complexity:
 - We plan to complete testing using latest FDR-2 data samples
 - We will test all Tier-1 sites using access to the DCS data
 - *Expect higher load from new DCS tests at all Tier-1*

