

# ATLAS Computing Status

---

Borut Kersevan  
Richard Mount

LHCC Referees September 2013

# ATLAS Resource Utilization in 2013




## ATLAS RESOURCE USAGE IN FIRST HALF OF 2013 (RRB)


	Location	Requested	Used
<b>CPU [kHS06]</b>	CERN	111	111
	Tier-1	316	435
	Tier-2	360	713
<b>Disk [PB]</b>	CERN	9.6	8.9
	Tier-1	35 [38]	35
	Tier-2	51 [52]	48
<b>Tape [PB]</b>	CERN	25	29 (incl. 9 PB of ESD)
	Tier-1	42	33

- ATLAS has utilized the computing resources in its Tiers well in the last year: *many thanks to sites for resources and excellent operating!*
  - We manage to provide a timely throughput of analyses to meet the physics requirements.
- An ongoing effort in software development to optimise the resource utilization by reducing the CPU consumption, event sizes - for Run-2...

Tiers CPU / running jobs



Tiers CPU / pending jobs



# Tier-1 Status



**WE HAVE MORE DISK AND CPU THAN REQUESTED,  
MANY THANKS TO TIER-1s!**

<i>Tier-1 Disk [PB]</i>	<i>Requested</i>	<i>Used</i>
RAW data	0.5	0.4
Real ESD+AOD+DPD	9	9
Simulated data	9	12
Calibration and alignment outputs	0	0
Group+User data	12	8
Cosmic ray data	0.2	0.2
Processing, scratch and I/O buffers	4.4	5.6
<b>Total</b>	<b>35 [38]</b>	<b>35</b>

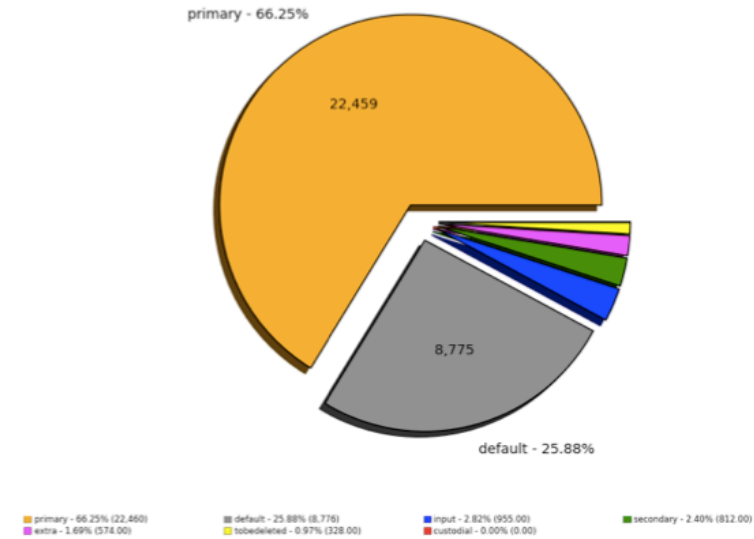
<i>Tier-1 CPU [kHS06]</i>	<i>Requested</i>	<i>Used</i>
Reprocessing	8	3
Simulation Production	80	190
Simulation Reconstruction	108	94
Group+User activities	120	148
<b>Total</b>	<b>316</b>	<b>435</b>

<i>Tier-1 Tape [PB]</i>	<i>Requested</i>	<i>Used</i>
RAW data	11	11
Real ESD+AOD+DPD		
Cosmics and other data	4	4
Simulated HITS+AOD	24	18
Group+User data	4	0.3
<b>Total</b>	<b>42</b>	<b>33</b>


**STILL, DISKS ARE FULL - WE ARE TAKING ACTIONS**



Number of Physical Bytes (in TBs) for 2013-08-27 (Sum: 33,906)



WallClock HEPSPROC6 Hours  
180 Days from Week 08 of 2013 to Week 34 of 2013



**CLEANUP DONE!  
(FastSim HITS)**

# Tier-2 Status



DISKS ARE ALSO FULL - WE ARE TAKING ACTIONS


WE HAVE MORE DISK AND CPU THAN PLEDGED,  
MANY THANKS TO TIER-2s!

<i>Tier-2 Disk [PB]</i>	<i>Requested</i>	<i>Used</i>
Real AOD+DPD	7	8
Simulated data	21	19
Calibration and alignment output	0.2	0.2
Group+User data	20	13
Processing, scratch and I/O buffers	3.4	7.2
<b>Total</b>	<b>51 [52]</b>	<b>48</b>

<i>Tier-2 CPU [kHS06]</i>	<i>Requested</i>	<i>Used</i>
Simulation Production	158	308
Simulation Reconstruction	42	207
Group + User activities	160	198
<b>Total</b>	<b>360</b>	<b>713</b>



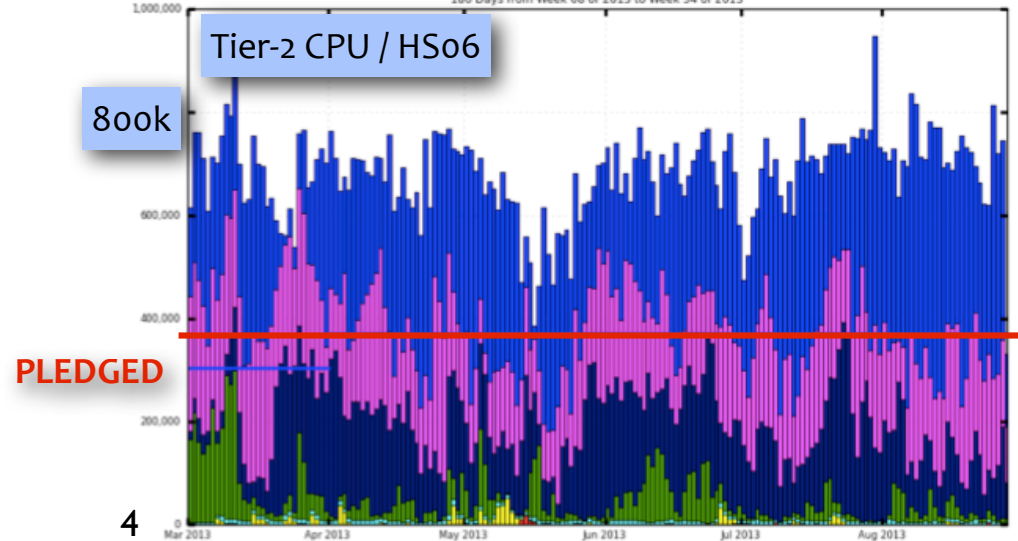
Number of Physical Bytes (in TBs) for 2013-08-28 (Sum: 39,081)  
primary - 56.80%



primary - 56.80% (22,197)    default - 26.54% (10,373)    secondary - 12.56% (4,907)    input - 2.90% (1,133)  
tobedeleted - 0.92% (360.00)    extra - 0.29% (112.00)    custodial - 0.00% (0.00)



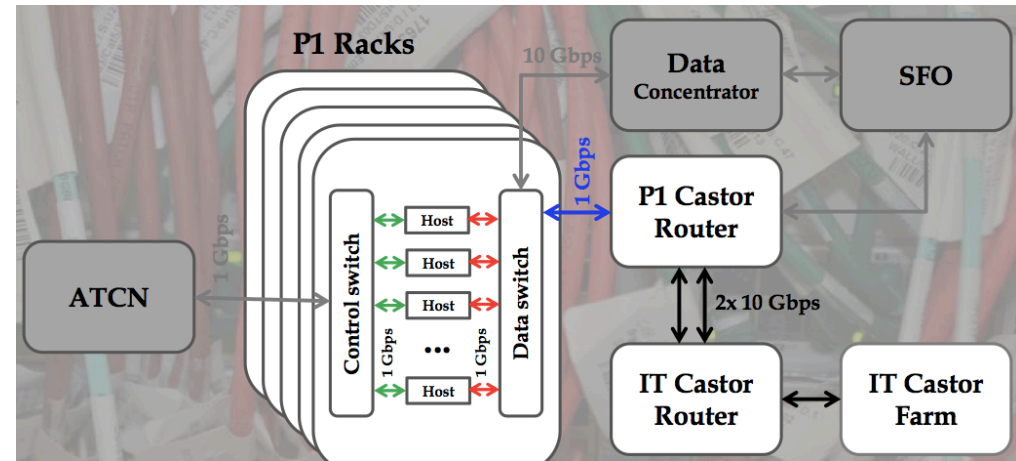
WallClock HEPSPC06 Hours  
180 Days from Week 08 of 2013 to Week 34 of 2013



# Highlight: Using the HLT Farm during LS1



- The computing experts from TDAQ, IT/SDC and BNL/T1 have set up the ATLAS HLT Farm to be used as a Grid 'site' during the LS1 as an **'opportunistic resource'**:
  - ~ 15,000 jobs = a big Tier-2!
  - A requirement from the C-RSG (RRB).
  - over the period of July-August: **140kHS06 28.3 billion seconds of wall clock time, job efficiency of 94.7%, completing 815k jobs.**
  - **A big success! Many thanks to all involved parties!**
  - The idea of using Cloud middleware (OpenStack) as the overlay infrastructure:
    - CERN IT (Agile), CMS (HLT Farm) and BNL all on OpenStack:
      - Similar use cases:
        - support if needed,
        - sharing experiences,
        - BNL has already part of its resources 'cloudified' and ATLAS is successfully using them!
      - An excellent use-case to gain experience with Cloud technologies!



# Data Placement Policies And Disk Cleanup

---



- Replication and data placement for Run-2: **somewhat more aggressive (!)**, the plan is to start implementing them **now** to optimize disk space usage!
  - **The number of pre-placed (primary) replicas of AODs will go down to 1 replica at Tier-1 and 1 replica at Tier-2 (almost done).**
    - We will rely even more on popularity-based dynamic data placement (PD2P) and automatic cleanup of secondary copies (Victor).
    - Consequently, more emphasis on the network use is expected.
  - **We will introduce an automated algorithm to clean (set to secondary) the pre-placed replicas in centrally managed disks, for example:**
    - If a data set was not accessed for six months set the Tier-1 replica as secondary.
    - If a data set was not accessed for one year set the Tier-2 replica as secondary.
    - The parameters will need operational fine-tuning.
    - **As a consequence example, we will use tape retrieval from Tier-1s more :**
      - e.g. the unpopular AODs will remain only on tape, retrieved for derivation of group-specific data or (more rarely) user analysis..

# Why the Change in Policy: Current Snapshot of Unused Data



- **The T1 and T2 datadisks are full of primary replicas, it is impacting our production!** (tasks not assigned to tiers with input data but no place for outputs...).
- (On-going) Clean-up action of removing all the datasets that were not touched since January 2013.
  - The total volume of untouched data is ~4PB in T1 and ~4PB in T2s (about 9% of total)

