



ATLAS Computing Resources

Dario Barberis

CERN & Genoa University/INFN



Computing Model and Resources

- The ESD size kept increasing during the last few years, as the ability to (partially) reprocess the events starting from the ESD needs more input information (calorimeter cells, InDet hits)
 - We think that, as we understand the detector better, we will be able to reduce the amount of extra information in the ESD
- We have therefore to provide for larger ESD in early years, reaching design value for higher luminosity in 2010:
 - Still can only allow ~1 MB/event in 2007/8 (instead of 0.5 MB/event original design)
- Having now a much better detector geometry representation than in the past, and operating the default simulation over $|\eta| < 6$ (instead of $|\eta| < 3$), the simulation time per event increased from 100 kSI2k-sec/event to 400 kSI2k-sec/event on average
 - We think that for many channels we will be able to use shower parameterisation in the calorimeters, but its performance is still under test
- In the first years of operation, there will be the need to tune calibration and reconstruction algorithms on real data; we have therefore increased the available CPU for user reconstruction with a decreasing profile with time
- The new resource calculation takes into account the agreed machine schedule for the first years of operation, the above changes in ATLAS input numbers and the global envelope of resource pledges according to the WLCG MoU



LHC schedule used for resource calculations

<i>year</i>	<i>energy</i>	<i>luminosity</i>	<i>physics beam time</i>
2007	450+450 GeV	5×10^{30}	protons - 26 days at 30% overall efficiency → 0.7×10^6 seconds
2008	7+7 TeV	0.5×10^{33}	protons - starting beginning July → 4×10^6 seconds ions - end of run - 5 days at 50% overall efficiency → 0.2×10^6 seconds
2009	7+7 TeV	1×10^{33}	protons: 50% better than 2008 → 6×10^6 seconds ions: 20 days of beam at 50% efficiency → 10^6 seconds
2010	7+7 TeV	1×10^{34}	TDR targets: protons: → 10^7 seconds ions: → 10^6 seconds




Total ATLAS Requirements start 2008

	CPU (MSI2k)	Disk (PB)	Tape (PB)
Tier-0	3.7	0.15	2.4
CAF	2.1	1.0	0.4
Sum of Tier-1s	18.1	9.9	7.7
Sum of Tier-2s	17.5	7.7	0.0




Evolution

New T0 Evolution




	2007	2008	2009	2010	2011	2012
◆ Total Disk (TB)	75.14785714	152.4621429	277.3242857	472.3528571	472.3528571	472.3528571
■ Total Tape (TB)	381.3075	2381.711	5267.2345	11371.158	17475.0815	23579.005
▲ Total CPU (kSI2k)	1910	3705	4058	6105	6105	6105

New T1 Evolution




	2007	2008	2009	2010	2011	2012
◆ Total Disk (TB)	2157.0332	9938.696929	19686.41793	39487.79764	56190.82307	72893.8485
■ Total Tape (TB)	1543.186667	7693.996427	14949.57676	28698.0172	44929.67775	63644.55841
▲ Total CPU (kSI2k)	3173.323529	18122.83529	28423.02353	49573.22353	70723.42353	91873.62353

New CAF Evolution



	2007	2008	2009	2010	2011	2012
◆ Total Disk (TB)	212.2436607	986.3915464	1529.026057	2777.498914	4047.976771	5255.197486
■ Total Tape (TB)	57.3206625	356.5720482	625.1016482	1017.151648	1342.801648	1668.451648
▲ Total CPU (kSI2k)	821	2069	2502	4596	6523	8450


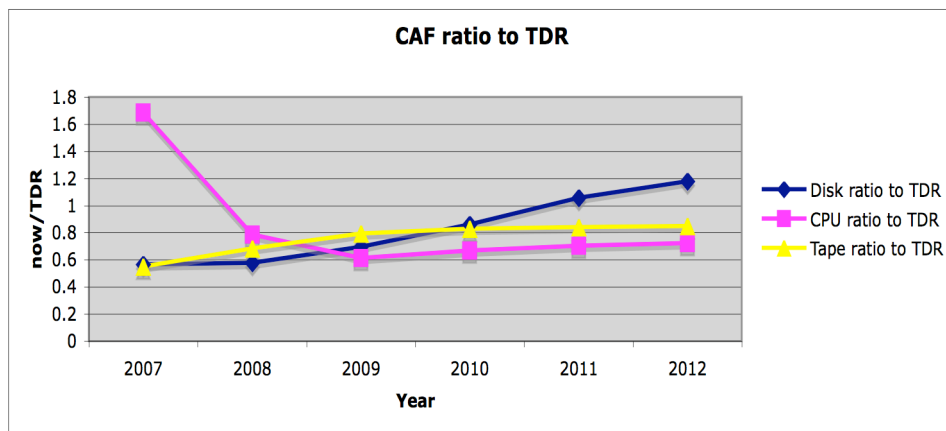


New T2 Evolution



	2007	2008	2009	2010	2011	2012
◆ Disk (TB)	1259.040486	7744.368955	13112.03563	22132.30423	31091.45139	40050.91999
■ CPU (kSI2k)	2336.108333	17494.50644	26972.75589	51544.63737	69128.41886	86712.20034



Ratio to Computing TDR



NB: there was a mistake in the Tier-0/CAF CPU requirements for 2007 in the Computing TDR (subsequently corrected already in 2005): CPU for calibrations does not scale with the length of the data-taking period



Observations on computing resources

- Data storage requirements generally fall with reduced live-time (obviously)
- CPU does not fall as much
 - CERN CPU determined by rate and calibration requirements
 - More calibration and optimisation is needed for 2007 data
- Higher than hoped simulation time per event
- Tier-1s see significant reductions
 - Cumulative effect of less data on reprocessing
- Tier-2s see a small initial fall but are bigger after 2009
- There is an argument for spreading the gain and the pain with Tier-1s by introducing more flexibility in the model:
 - Tier-1s can now produce simulated data when not fully busy with reprocessing