

ATLAS DA JOB DURATION

Johannes Elmsheuser

Ludwig-Maximilians-Universität München, Germany

30 November 2010

ATLAS S&C Week, DA meeting

Update to presentation from 17 Nov 2010




Disclaimer:



- All numbers from 'Historical View Dashboard' at:
<http://dashb-atlas-job-dev.cern.ch/dashboard/request.py/dailysummary>
- Number cross checks look fine
- „Processed Events” only meaningful for athena jobs - no infos for ROOT jobs
- Some further cross checks for job „Walltime” necessary

ATLAS JOB STATISTICS I

ATLAS Grid jobs last month:




ATLAS Analysis jobs last month:




Large fraction of "terminated jobs" are DA jobs

ATLAS JOB STATISTICS II

ATLAS analysis CPU consumption last month:



ATLAS Analysis jobs Walltime last month:



- 250000 DA jobs per day consuming \approx 12000 CPU days per day
- \rightarrow 1.15 CPU hours per analysis job on average worldwide
- Assumption only Athena jobs with 200KB/Event: 16.9 Hz Event-rate
- (Same calculation leads to: 1.25 CPU days per production job)

ATLAS DA JOB DURATION I

More site details (my personal bias):

Sites	Jobs (last 30d)	Walltime d	Walltime h/job	Events/job	mode
BNL	945814	31527	0.29	3884	copy
Glasgow	76017	2534	1.06	9143	copy
DESY-ZN	107862	3595	1.07	11925	direct
FZK	318409	10613	1.09	5515	direct
AGLT2	579356	19311	1.15	3068	copy
DESY-HH	199833	6661	1.21	1795	direct
SLAC	761792	25393	1.41	4850	direct
LAPP	80239	2675	1.50	7858	copy
ROMA1	63710	2124	1.70	7469	copy

Several biases - so numbers might not be useful

- No distinction between input format mixture (ESD, AODs, D3PDs)
- No information about Events processed in ROOT jobs
- Some variations in average walltime hours/job, dependent on SE architecture and CPU speed

ATLAS DA JOB DURATION II

- Still some sites see many short DA jobs
- pathena job splitting controlled via schedconfig parameter `maxinputsize` for copy-to-scratch sites or 50 files/subjob if not overruled by user
- GangaPanda splits so far via:
`config.Athena.MaxFileSizePandaDQ2JobSplitter=13336 MB` - has been changed in Ganga 5.5.18 to 50 files or `maxinputsize`
- Script lists Panda schedconfig parameter `maxinputsize`:
`/afs/cern.ch/user/e/elsheus/public/queues.sh`
- More direct access sites ? leave it up to site/cloud to decide (?)

OVERVIEW (I) - SITE/MAXINPUTSIZE/DIRECT

ANALY_AGLT2	14336	False	ANALY_GLASGOW_GLEXEC	14336	False	ANALY_LRZ	14336	True
ANALY_ALBERTA-WG1	36000		ANALY_GLOW-ATLAS	14336	False	ANALY_LYON-T2	30000	False
ANALY_ANLASC	14336		ANALY_GOEGRID	14336	True	ANALY_LYON_DCACHE	30000	False
ANALY_ARC	204800		ANALY_GRIF-IRFU	14336		ANALY_MANC	14336	False
ANALY_AUSTRALIA	14336	False	ANALY_GRIF-LAL	14336		ANALY_MPPMU	14336	True
ANALY_BEIJING	14336		ANALY_GRIF-LPNHE	14336		ANALY_MWT2	14336	True
ANALY_BHAM	14336	False	ANALY_HEPHY-UIBK	14336	True	ANALY_MWT2_X	14336	True
ANALY_BNL-ATLAS_1	50000		ANALY_HU-ATLAS_Tier2	14336		ANALY_NCG-INGRID-PT	14336	
ANALY_BNL_GLEXEC	50000		ANALY_IFAE	14336		ANALY_NET2	14336	False
ANALY_BNL_LOCAL	50000		ANALY_IFIC	14336		ANALY_NIKHEF-ELPROD	14336	
ANALY_BNL_test	50000		ANALY_IHEP	14336		ANALY_OU_OCHEP_SWT2	14336	
ANALY_BNL_test2	50000		ANALY_IL-TAU-HEP	14336		ANALY_OX	14336	False
ANALY_BNL_test3	50000		ANALY_INFN-FRASCATI	50000		ANALY_PIC	50000	True
ANALY_BRANDEIS	14336		ANALY_INFN-MILANO-ATLAS	14336		ANALY_PNPI	14336	
ANALY_CAM	14336	False	ANALY_INFN-NAPOLI	14336		ANALY_QMUL	14336	
ANALY_CERN	14336	False	ANALY_INFN-ROMA1	14336		ANALY_QMUL_TEST	14336	
ANALY_CERN_XROOTD	14336	False	ANALY_INFN-T1	50000		ANALY_RAL	50000	
ANALY_CPPM	14336		ANALY_ITEP	14336		ANALY_RALPP	14336	False
ANALY_CSCS	14336	True	ANALY_IllinoisHEP	14336		ANALY_RAL_XROOTD	50000	
ANALY_CSTCDIE	14336		ANALY_JINR	14336		ANALY_RHUL	35000	False
ANALY_CYF	14336	False	ANALY_LANCS	14336	False	ANALY_ROMANIA02	14336	
ANALY_DESY-HH	14336	True	ANALY_LAPP	14336		ANALY_ROMANIA07	14336	
ANALY_DESY-ZN	14336	True	ANALY_LIP-Coimbra	14336		ANALY_RRC-KI	14336	
ANALY_DRESDEN	14336		ANALY_LIP-Lisbon	14336		ANALY_SARA	50000	
ANALY_DUKE	14336		ANALY_LIV	14336	False	ANALY_SCINET	14336	
ANALY_ECDF	14336	False	ANALY_LONG_BNL-ATLAS	50000		ANALY_SFU	9000	
ANALY_FREIBURG	14336	True	ANALY_LONG_BNL_LOCAL	50000		ANALY_SFU_bugaboo	21000	
ANALY_FZK	50000	True	ANALY_LONG_LYON-T2	30000	False	ANALY_SHEF	14336	False
ANALY_FZK_CP	50000	False	ANALY_LONG_LYON_DCACHE	30000	False	ANALY_SLAC	14336	True
ANALY_FZU	14336	False	ANALY_LPC	14336		ANALY_SLAC_LMEM	14336	True
ANALY_GLASGOW	14336	False	ANALY_LPSC	14336		ANALY_SWT2-CPB	14336	True

OVERVIEW (II) - SITE/MAXINPUTSIZE/DIRECT

ANALY_TAIWAN	14336	
ANALY_TECHNION-HEP	14336	
ANALY_TEST	25000	False
ANALY_TOKYO	14336	
ANALY_TR-10-ULAKBIM	14336	
ANALY_TRIUMF	25000	False
ANALY_TW-FTT	14336	
ANALY_Tufts_ATLAS_Tier3	14336	False
ANALY_UAM	14336	
ANALY_UCL	18000	False
ANALY_UNI-DORTMUND	14336	
ANALY_UTA_T3	14336	
ANALY_VICTORIA	14336	
ANALY_VICTORIA-WG1	23000	
ANALY_WEIZMANN	14336	
ANALY_WISC-ATLAS	14336	False
ANALY_wuppertalprod	14336	True