

ATLAS FALL2010 REPROCESSING @ CNAF

- Atlas point of view
 - metrics
 - job workflow
- CNAF point of view:
 - tape recall
 - storage usage
 - GPFS network

CNAF reprocesses 5% of total p - p data taken during 2010

Reprocessing in two phases:

- Phase I: data collected until September 2010 (42 TB @ CNAF)
 - Jobs sent in three bulks
 - Done from Oct. 30th to Nov 10th
 - 23745 jobs
- Phase II: data collected during October 2010 (12 TB @ CNAF)
 - Done from Nov. 17th to 24th
 - 5472 jobs

Additional output merging-jobs (17000) run from Oct. 30th to Nov 20th

Metric: reprocess 100% of the assigned jobs in the given time interval (3 weeks)

Cnaf behaved well: average efficiency comparable to the other atlas T1s

- ~3000 jobs running daily
- Bad point: more than 6000 failures! (20%) Mainly due to temporary down of SRM/LFC
- Good point: problems always fixed in a reasonable time.

October-November 2010

Generated on Thu Nov 18 17:10:59 2010 (CERN) [[previous reports](#)]

All ATLAS T1s reprocessing stats

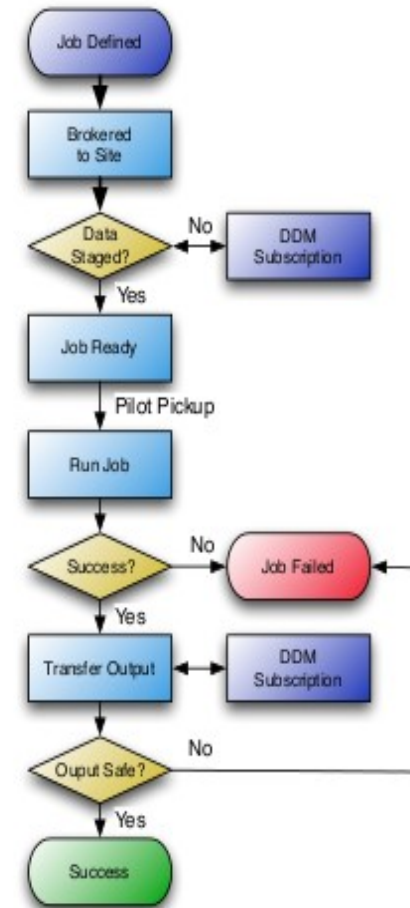
Long period without running jobs:
With the available CPU's, CNAF
can run even more jobs

date-tier	CA	CERN	DE	ES	FR	IT	ND	NL	TW	UK	US	sum
10-10-30	88	0	974	976	908	1047	945	1934	2199	3355	2760	15186
10-10-31	2541	0	1	212	579	77	1475	2272	85	3249	11006	21497
10-11-01	2530	0	3619	73	7420	4448	628	3737	1866	3024	8004	35349
10-11-02	2714	109	5649	954	4463	5771	1295	2154	2670	5720	7902	37401
10-11-03	4891	40	9952	4590	2554	27	1924	2903	2003	3250	14146	46280
10-11-04	7	0	3625	163	3289	17	1511	842	2225	1	15809	27409
10-11-05	1539	10	0	4043	4317	922	2836	2564	1927	9574	21384	49116
10-11-06	2989	0	3598	4598	10722	3309	2482	1029	2584	11235	4648	47194
10-11-07	6917	0	5722	3943	7191	2801	4073	0	2554	20794	14315	68310
10-11-08	6749	0	15858	3643	5807	7003	6103	4084	334	8655	15161	73397
10-11-09	24	3989	1432	3180	6752	319	219	7075	597	298	19584	43469
10-11-10	0	6587	1	4	3026	2	7	9304	1806	193	7777	28707
10-11-11	0	5517	0	3	58	0	1	14324	2722	9	24361	46995
10-11-12	0	2	0	0	239	0	0	6994	2077	0	13677	22989
10-11-13	0	0	0	0	7161	0	0	979	638	0	0	8781
10-11-14	0	0	0	0	4905	0	0	663	0	0	0	5569
10-11-15	0	0	0	0	9402	0	0	0	0	0	0	9402
10-11-16	0	0	0	0	63	0	0	0	0	0	0	63
10-11-17	0	0	0	0	1	0	0	0	0	0	0	1
total jobs	30992	16256	50434	26385	78788	23745	23503	60864	26289	69361	180544	587161
total done	30989	16254	50431	26382	78781	23743	23499	60858	26287	69357	180534	587115
%%	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
aborted	3	2	3	3	7	2	4	6	2	4	10	46
%%	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
running	0	0	0	0	0	0	0	0	0	0	0	0
%%	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

date-tier	CA	CERN	DE	ES	FR	IT	ND	NL	TW	UK	US	sum
10-11-17	1673	512	1834	1729	0	953	342	282	276	1087	3153	11841
10-11-18	1959	4808	2620	2546	2659	1727	1829	3285	1855	945	8331	32564
10-11-19	66	3013	4839	1607	1361	2174	2830	3766	1044	17	9426	30143
10-11-20	9	5287	3620	318	192	538	419	4225	1696	2	7931	24237
10-11-21	105	100	259	2	3	26	20	926	436	0	2618	4495
10-11-22	0	17	44	2	1	6	14	39	0	0	40	163
10-11-23	346	223	1	0	0	41	1	1	0	270	8	891
10-11-24	18	3816	0	0	0	6	1	0	0	2982	0	6823
10-11-25	22	53	0	0	0	0	0	0	0	73	0	148
total jobs	4198	17838	13217	6204	4216	5471	5457	12524	5307	5395	31509	111336
total done	4198	17829	13217	6204	4216	5471	5456	12524	5307	5376	31507	111305
%%	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.7	100.0	100.0
aborted	0	0	0	0	0	0	1	0	0	0	1	2
%%	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
running	0	9	0	0	0	0	0	0	0	19	1	29
%%	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.0	0.1

Typical ATLAS job workflow

<i>Panda Job State</i>	<i>Description</i>
Defined	Job defined by production system
Assigned	Job brokered to a site
Activated	Job's data ready at site
Running	Job executing
Holding	Job outputs being registered in DDM
Transferring	Job outputs awaiting transfer to TI
Finished	Job completed successfully
Failed	Job completed unsuccessfully




Tape recall of RAW data is performed when jobs are in **assigned** status. If data not staged, a DDM subscription starts (using DQ2)

Each job recalls and reprocesses a unique file (average size 2GB)


As jobs are sent in bulks, many data are prestaged at the same time, before the jobs start to run

Tape recall 1

Atlas panda monitor



Cyan line: assigned jobs



Red line: data copied from tape to disk buffer

(notice bulk submission structures)

CNAF TSM network monitor

Tape recall 2

From TAPE network monitor and logs:

- Phase I: 26061 files recalled (~50 TB)
 - BUT only 20182 files have been recalled AND reprocessed
 - 5483 files recalled twice; 376 files recalled > 2 times
 - Time elapsed between two further recall : 20 hours
 - Time consistent with ATLAS job re-submission time after a failure

Files flushed out of buffer because of short pinning time:
increased buffer helped to fix the problem
- Phase II: 11276 files recalled (~10.5 TB)
 - No double recall
 - Affected by network down on 19th Nov.
- One more (apparent) inconsistency:
ATLAS jobs are 23745, BUT only 20182 have been staged and processed:
perhaps the remaining ~3300 files already resident on buffer since
export from T0

Tape recall 3

TSM disk buffer traffic: Nov. 1st -2nd


10,6 TB recalled

Green: buffer-->tape throughput (p-p and HI T0 export)

Red: tape--> buffer throughput (RAW data to be reprocessed)

10.6 TB recalled

GPFS utilization




Peaked structure because of the low performing buffer (60 TB shared with MCTAPE area)

Buffer migrated on a more performant DS and size increased to 117TB


Number of file recalled for each tapeserver

TSM recall OK



Tape recall 4

GPFS utilization




Smoother file recall after buffer size increased

Some failures observed on Nov 5th


39.6 TB recalled

During these days multiple recalls observed

TSM recall OK



TSM recall FAILED




Tape recall 5


Phase II: Nov 17th -19th

10 TB recalled, 2 TB already on buffer

GPFS utilization




TSM recall OK




Summary last month


GPFS utilization




TSM recall OK



Storage occupancy





Disk and tape occupancy consistent with 2010 pledges





Network utilization gpfs --> farm

One day of user analysis



One week of user analysis



User analysis jobs


Running jobs (dark green) and traffic on gpfs have similar trend (low traffic when few jobs running)

User jobs access ESD data (large size) → network saturation at 4.6 Gb/s

In the daily plot, NO reprocessing running

Network utilization gpfs --> farm

Network utilization



Since July 1st 2010 the load of traffic is increased owing to user analysis jobs

- Change networking configuration (increase actual limit to 10 Gb/s)
- Tuning user analysis job share (now: 40-60 user-production)

Reprocessing jobs DO NOT saturate the GPFS network

Conclusions

Reprocessing of atlas 2010 data run well

- All jobs terminated with good time efficiency
- Result can be even improved by keeping the services stable (SRM/LFC failures ~20%)
- Storage occupancy within 2010 pledges
- Repro jobs are not a problem for the networking (new setup will improve it)

Future: increase reprocessing load @ cnaf (10%?)

Backup

Problems:

- 4 Nov.: 18 jobs in pending status. Problem at atlas SS level. Simone fixed it
- 9 Nov.: GPFS stuck due to FS still mounted when VM shutdown
- 10 Nov. : LFC number of file limit reached. Limit increased
- 10 Nov. : jobs appended since pilot wrongly went to simul queue
- 11 nov again problem with LFC
- 11 nov problem with ce01
- 18 nov: srm down
- 20-21 nov: network down
- 21-23 Nov: BNL VOMS server missing