



ATLAS Software and Computing Week Status and plans

J. Shank

ATLAS S&C 4 April, 2011



Overview

- Highlights of the S&C week
- Computing resource usage in 2010
- Computing resource estimate for 2011-2013
- Upcoming meetings: TIM in Dubna, 31 May- 2 June



Some highlights for this S&C week

- Today's plenary
 - Guest talk about networking by Artur Barczyk.
- Tues.
 - ADC R&D and TaskForces
 - Open meeting with many invited from outside of ATLAS
 - Will have presentations and discussion on future computing projects relevant to ATLAS and other experiments.
 - International Computing Board meeting
 - Closed meeting. Note unusual time/day.
- Weds.
 - Extended session on CVMFS
 - Simulation. 2 sessions, 1 on Weds, plus extended session on Fri.
- Throughout the week, the usual sessions on ADC, PAT, DA, DDM, DB, Core SW, Group Production, PanDA, SIT, ...



2010 Computing Resource Usage

- ATLAS Computing Model
 - <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ComputingModel>
 - See [ATLAS computing usage report 2010.pdf](#)
 - And: [Resource Request for 2011-2013](#)
- Numbers presented in the Resource Request document above will be presented to the Resource Review Board on 12 April, 2011
 - It is important that you brief your funding agencies about our requests, particularly for 2012 where our requests have gone up because we now are running that year.



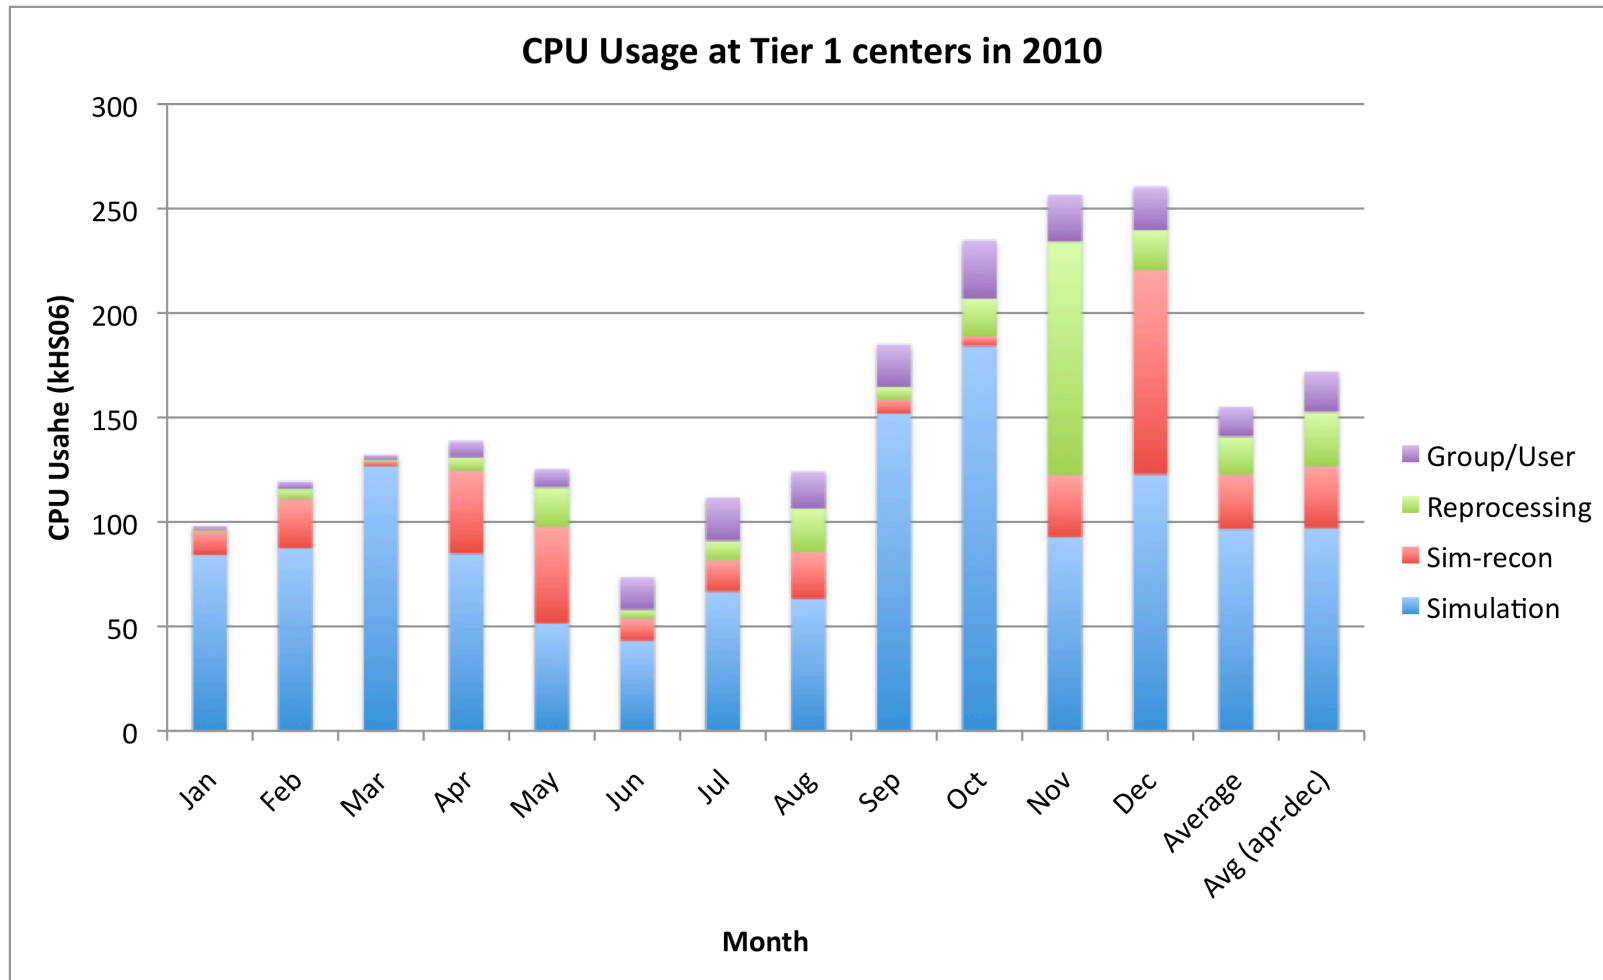
2010 Sizes and Reconstruction Times

- The "New 2010" column was used as the basis of estimates for last year.

LHC and data taking parameters		New 2010	Actual
Rate [Hz]	Hz	200	280
Time [sec]	MSeconds	7.3 (2009+2010)	3.6 (2010)
Real data	BEvents	1.46	1.0
Simulation	BEvents	0.6	1.4
Fast Simulation	BEvents	0.6	0
Event sizes			
Real RAW	MB	1.6	1.4
Real ESD	MB	0.8	1.6
Real AOD	MB	0.15	0.18
Sim RAW	MB	2	2
Sim ESD	MB	1.1	2
Sim AOD	MB	0.18	0.4
CPU times per event			
Full sim	HS06 sec	6000	4100
Real recon	HS06 sec	80	140
Sim recon	HS06sec	135	160



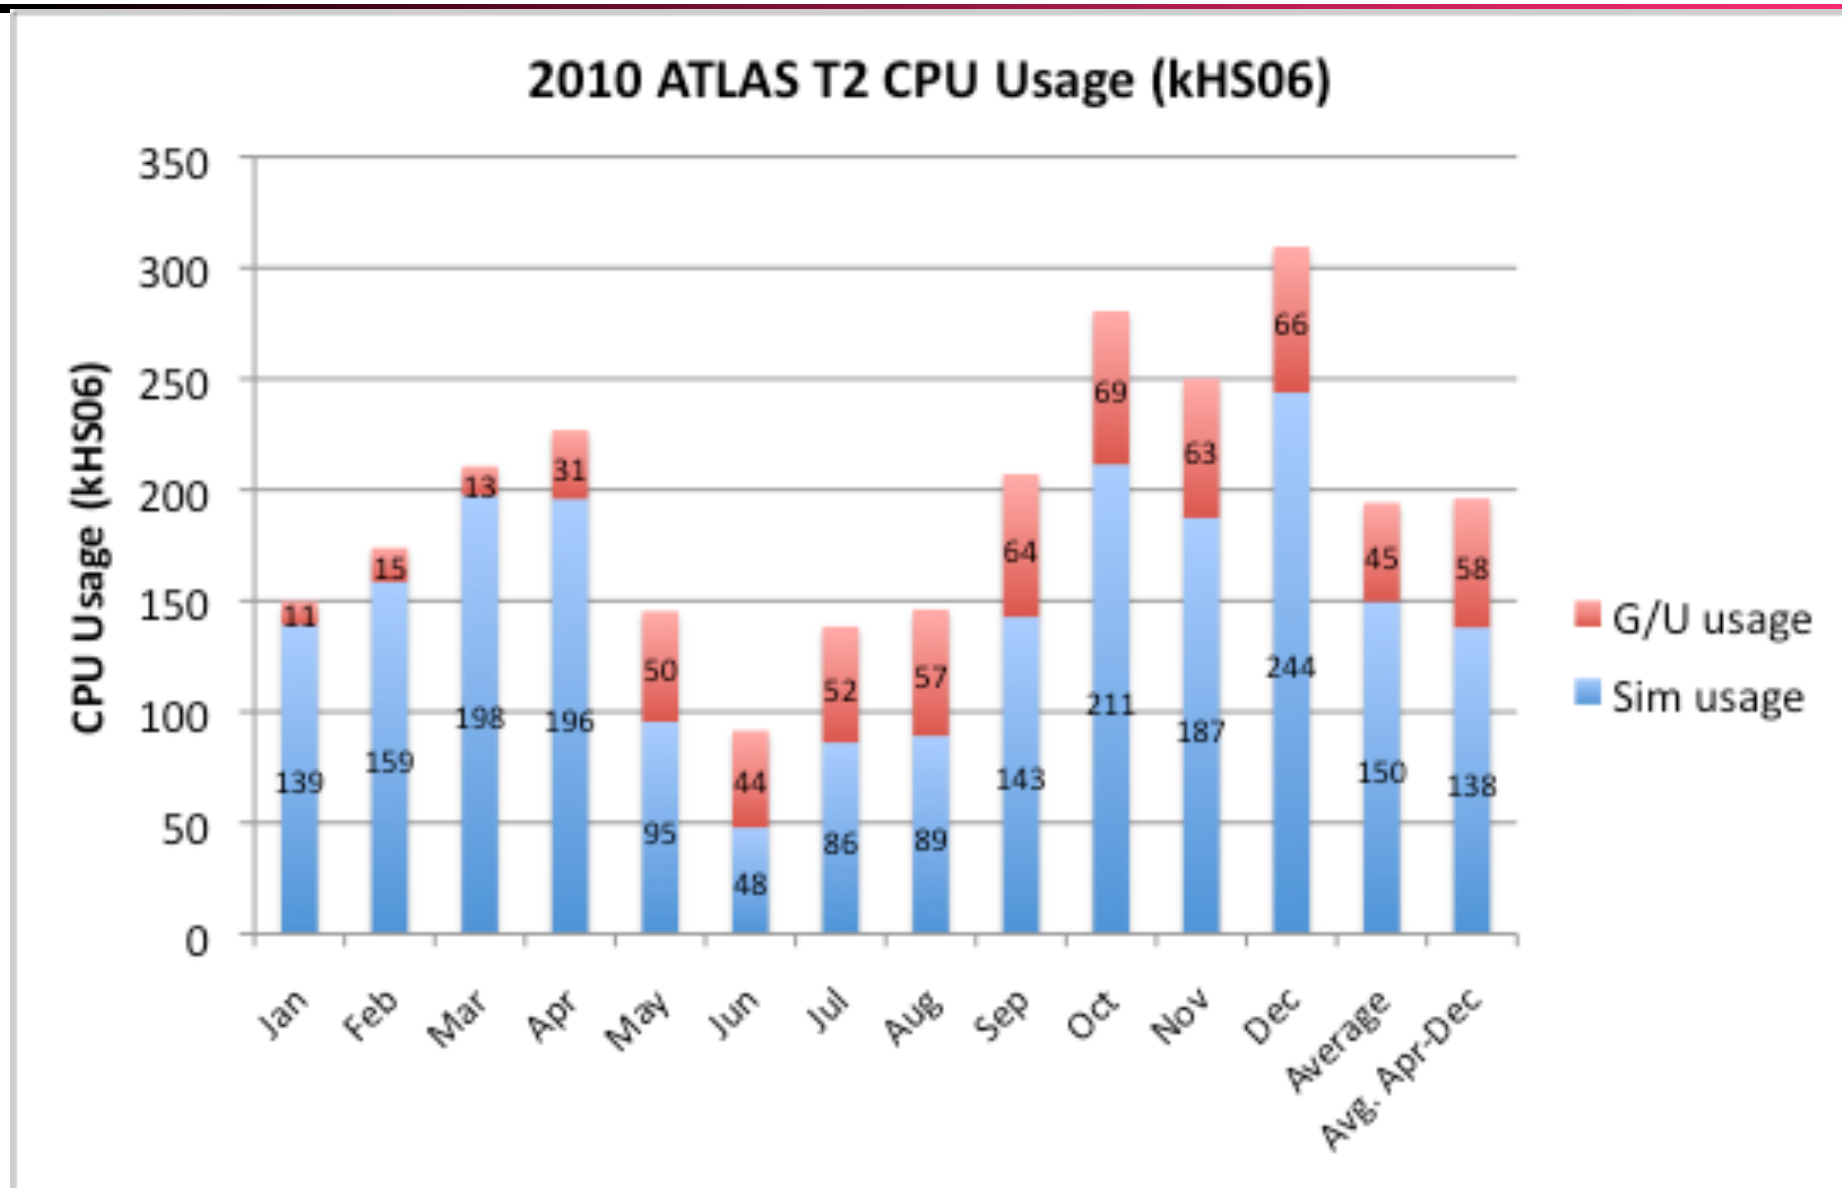
CPU Usage at Tier 1 in 2010



We are still working to improve accounting, which caused some confusion between categories



CPU Usage at Tier 2 in 2010





2010 Usage summary

- This was an initial data taking year. In some ways, not reflecting usage patterns of steady high-luminosity running.
 - Machine live-time was low and weighted towards the end of year
 - Data were collected with rapidly evolving LHC beam conditions
 - Understanding of the experimental performance was also evolving rapidly
 - Analysis in the first $\frac{1}{2}$ year more detector commissioning and performance focused
 - Changing rapidly at the end of 2010
 - Simulation samples focused on low initial luminosity and early physics.
 - Group and User CPU activity less than predicted but increasing at the end of the year



Full Model Spreadsheet

- To make computing resource estimate, we use a large spreadsheet that has evolved over the years and which partly reflects how we actually used resources
- Basic assumptions:

	Rate(Hz)	sec/year	Events/y	Size(MB)
RAW	200	5.18E+06	1.04E+09	2.8
ESD	200	5.18E+06	1.04E+09	2.2
AOD	200	5.18E+06	1.04E+09	0.36
TAG	200	5.18E+06	1.04E+09	0.001

Pile up x2

- Reconstruction time: 200 HS-sec
- Baseline Trigger rate: 200 Hz
 - Optimize rate to maximize physics output within the existing computing resources. We would benefit from a higher rate than 200Hz.



The 2011-12 Data Distribution Model

- The model:
 - 1 Copy of RAW data on disk distributed over all T1s
 - RAW event is 1.4 doubled now to 2.8 for pileup
 - A rolling buffer of 10% of ESD distributed over T1s
 - ESD size reduced to 1.1 MB/evt but doubled now due to pileup → 2.2 MB
 - ESD from "small streams" (equivalent to 10% ESD) distributed to T1s
 - 10 copies of AOD distributed over all clouds
 - Only 3 copies of previous version AOD
 - 2 copies of DESD at T1, 4 at T2.
 - Sum of all DESD = size of AOD
- Other changes that affect Computing Resources:
 - Reprocessing: 100HS-sec (from 2010 data) assumes $P_{\tau} > 400$ MeV instead of $P_{\tau} > 100$ MeV
 - Then doubled due to pileup → 200HS06-sec
 - In tables on next slides:
 - Green: new request is within 5% of 2010 RRB request
 - Red: new request is more than 2010 RRB request by more than 5%
 - Violet: new request is less than 2010 RRB request by more than 5%



Rolled up Summary of Request

CPU [kHS06]	2010	2011	2012	2013
CERN	74	74	73	60
Tier-1	178	202	259	280
Tier-2	226	275	295	321
Disk [PB]				
CERN	4.7	7	9	10
Tier-1	22	22	27	30
Tier-2	24	35	49	56
Tape [PB]				
CERN	9	14	18	18
Tier-1	18	28	36	40



CERN CPU

CERN CPU (kHS06)	2010	2011	2012	2013
CERN CPU Total	74	74	73	60
Tier-0 subtotal	30	38	38	33
T0: Full reconstruction	16	28	28	28
T0: Partial processing and validation	3	2	2	2
T0: Merging and monitoring	3	2	2	2
T0: Automatic calibration	7	5	5	0
T0: Servers	1	1	1	1
CAF subtotal	44	36	35	27
CAF: Partial, reconstruction, debugging and monitoring	3	6	4	0
CAF: Non-automatic calibrations	4	4	4	0
CAF: Group activities	22	11	12	12
CAF: User activities	3	3	4	4
CAF: Servers	12	12	12	12

Main effect:

- Running in 2012
- CPU shifted from CAF to Tier0 to cope with pile-up

Requests presented to RRB in 2010

IF CPU (kHS06) RRB Oct., 2010	2010	2011	2012	2013
Partial, reconstruction, debugging and monitoring	3	3	1	3
Non-automatic calibrations	4	4	0	4
Group activities	22	22	19	16
User activities	3	5	3	5
Servers	12	12	12	12
Total	44	46	36	40
Tier-0 CPU (kHS06) RRB Oct., 2010	2010	2011	2012	2013
Full reconstruction	16	16	2	16
Partial processing and validation	3	3	0	3
Merging and monitoring	4	4	2	4
Automatic calibration	7	7	7	7
Servers	1	1	1	1
Total	31	31	11	31

Total CERN CPU (2010 RRB) 75 76 47 71



CERN Disk

<i>Tier-0 Disk (PB)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
<i>CERN Disk Total</i>	4.7	7	9	10
<i>Tier-0 subtotal</i>	0.85	0.85	0.85	0.85
T0: Buffer for RAW and processed data	0.65	0.65	0.65	0.65
T0: Buffers for merging	0.10	0.10	0.10	0.10
T0: Tape buffer	0.10	0.10	0.10	0.10
<i>CAF subtotal</i>	3.8	6.3	8.5	9.3
CAF: Calibration and alignment	0.6	0.7	0.7	0.7
CAF: Derived detector data	1.1	2.6	4.1	4.1
CAF: Derived simulated data	0.8	1.5	2.3	3.1
CAF: Group data (grid aware)	0.9	1.0	1.0	1.0
CAF: Group data (non grid aware)				
CAF: User data (grid aware)	0.4	0.4	0.4	0.4
CAF: User data (non grid aware)				

Requests presented to RRB in 2010

<i>CAF Disk (PB)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
Calibration	0.6	0.6	0.6	0.7
Derived data	1.1	2.0	2.0	2.9
Derived simulated	0.8	1.2	1.7	2.2
Group data	0.9	1.9	1.9	1.2
Group data (non grid aware)				
User data	0.4	0.4	0.4	0.0
User data (non grid aware)				
<i>CAF Total</i>	3.8	6.1	6.6	7.0
<i>T0 Total</i>	0.9	0.9	0.9	0.9
<i>CERN Total</i>	4.7	6.9	7.4	7.9



Tape at CERN

<i>Tier-0 Tape (PB)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
<i>Total</i>	8.9	14	18	18

Requests
presented to
RRB in 2010



Increase due to raw event size
with pile-up

<i>Tier-0 Tape</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
<i>Total</i>	8.9	12.2	12.5	15.7



T1 Resources - CPU

Tier-1 CPU (kHS06)	2010	2011	2012	2013
Re-processing	15	41	71	74
Simulation production	109	77	77	77
Simulation reconstruction	16	28	39	52
Group (+user) activities	38	56	72	76
Total	178	202	259	280

Tier-1 CPU
(kHS06) RRB Oct., 2010

	2010	2011	2012	2013
Re-processing	34	34	24	70
Simulation production	108	109	108	146
Simulation reconstruction	3	34	35	47
Group (+user) activities	33	49	57	77
Total	178	226	225	340

← Requests presented to RRB in 2010



T1 Resources - Disk

<i>Tier-1 Disk (PB)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
Current RAW data	3.9	4.3	4.3	4.3
Real ESD+AOD+DPD data	6.5	3.8	5.4	5.4
Simulated RAW+ESD+AOD+DPD data	4.1	4.4	5.6	6.8
Calibration and alignment outputs	2.2	0.4	0.4	0.4
Group data	1.4	4.5	6.5	7.6
User data (scratch)	0.6	0.6	0.6	0.6
Cosmics	0.2	0.2	0.2	0.2
Processing and I/O buffers	3.0	3.6	4.3	4.4
Total	22	22	27	30

Requests presented to RRB in 2010

<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
3.9	0.7	0.7	1.0
6.5	11.0	11.0	15.8
3.9	6.1	7.7	9.4
1.8	1.0	0.0	0.4
2.3	2.3	2.6	4.3
0.6	0.6	0.6	0.6
0.2	0.2	0.2	0.2
3.0	3.0	3.5	6.5
22.2	24.9	26.4	38.3



T1 Resources - Tape

<i>Tier-1 Tape (PB)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
LHC RAW data	3	6	9	9
Real ESD+AOD+DPD data	3	6	7	8
Simulated data	8	11	14	17
Group + User	0	2	2	2
Cosmics and other data	4	4	4	4
Total	18	28	36	40

<i>Tier-1 Tape (PB) RRB Oct., 2010</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
LHC RAW data	3	5	5	7
Real ESD+AOD+DPD data	3	7	10	15
Simulated data	8	12	16	21
Group + User	0	2	4	6
Cosmics and other data	4	4	4	4
Total	18	30	39	53

← Requests presented to RRB in 2010



T2 Resources - CPU

<i>Tier-2 CPU (kHS06)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
Simulation production	65	56	56	56
Group activities	38	56	58	58
User activities	123	163	180	207
Total	226	275	295	321

<i>Tier-2 CPU (kHS06)</i>	<i>2010</i>	<i>2011</i>	<i>2012</i>	<i>2013</i>
Simulation production	65	65	65	89
Group activities	38	49	52	62
User activities	123	164	178	254
Total	226	279	295	405

Requests presented to RRB in 2010



T2 Resources - Disk

Tier-2 Disk (PB)	2010	2011	2012	2013
Current RAW data	0.4	0.0	0.0	0.0
Real AOD+DPD data	11	13	21	21
Simulated RAW+ESD+AOD+DRD data	8	12	16	21
Calibration and alignment output	0.0	0.3	0.3	0.3
Group data	3	7	9	10
User data (scratch)	1	2	2	3
Processing buffers	1	1	1	1
Total	24	35	49	56

Tier-2 Disk (PB) RRB Oct., 2010	2010	2011	2012	2013
Current RAW data	0.4	0.3	0.3	0.3
Real AOD+DPD data	11	20	20	29
Simulated RAW+ESD+AOD+DRD data	8	11	15	20
Calibration and alignment output	0.0	0.3	0.3	0.3
Group data	3	5	5	6
User data (scratch)	1	2	2	3
Processing buffers	1	1	1	1
Total	24	38	44	60

← Requests presented to RRB in 2010



Summary of Resource Estimates

- The change in LHC running (full running in 2012, shut-down in 2013) requires more resources in 2012 but generally below those previously envisioned for 2013
- Higher pileup increases reconstruction time, requiring more CPU at the T0, ~20% more, which will be redeployed from CAF.
- Higher pileup increases event sizes, that we cope with by a new data distribution model that phases out ESD and reduces the number of replicas of AOD and dESD.
- We are working hard to reduce further the impact of pile-up on reconstruction time and event sizes, but the largest gains are already achieved.

Conclusion: Planning of computing resources is a delicate balance between resources constraints (pledged in 2011, finances in 2012) and physics performance (trigger rate, reconstruction efficiency, data access, etc.). Our computing resource requirements are designed to optimize ATLAS physics output within available/predicted resources. We will continue to optimize operations as we gain further experience.



Database News

Database News

- 1) New version of Oracle (11g) will be introduced at the end of 2011
A test RAC is now available for application tests
All database application developers must get in contact with Gancho to schedule these tests during next few weeks

- 2) Next step of the consolidation of distributed databases
The 3D databases at NDGF, SARA and CNAF will be decommissioned next week
There should be no consequence on any application as Squids are active at all these sites, but just in case...



Upcoming Meeting

- ATLAS Computing and Technical Interchange Meeting
 - 31 May 2011 to 02 June 2011. Joint Institute For Nuclear Research (Dubna)

Tuesday 31 May 2011

Registration (08:00-09:00)

Welcome (09:00-09:30)

- Conveners: Korenkov, Vladimir; Prof. Rusakovich, Nikolai

Introduction and Russian Grid Segment (09:30-10:15)

- Conveners: Ilin, Viatcheslav; Korenkov, Vladimir

ATLAS Computing (10:15-10:45)

- Conveners: Shank, Jim

Tier-3 (11:00-12:30)

- Conveners: Benjamin, Doug; Filipcic, Andrej

Tier-3 (13:30-15:00)

- Conveners: Benjamin, Doug; Filipcic, Andrej; Andreeva, Julia

ATLAS Distributed Computing. PanDA, DDM, PD2P. (15:20-18:00)

- Conveners: De, Kaushik; Dr. Klimentov, Alexei

Welcome drink (18:30-20:00)



TIM Agenda

Wednesday 01 June 2011

RnD and TaskForces. Event level caching. Xrootd federation (08:30-11:30)

- Conveners: Fons Rademakers; Dr. Wenaus, Torre

RnD and TaskForces. Cloud Computing (11:30-14:00)

- Conveners: Campana, Simone; Panitkin, Sergey

River excursion and conference dinner (14:30-19:30)

Thursday 02 June 2011

RnD and Task Forces. Relational and non-relational databases. (08:30-11:30)

- Conveners: Canali, Luca; Dr. Garonne, Vincent (TBC)

Data Storage Evolution (11:30-12:30)

- Conveners: Duellmann, Dirk; Hanushevsky, Andrew

Data Storage Evolution (13:30-14:30)

- Conveners: Duellmann, Dirk; Hanushevsky, Andrew

Discussion and action items (14:30-16:30)

- Conveners: Dr. Klimentov, Alexei; Korenkov, Vladimir

Closing remarks (16:30-17:30)