

# Software news

**David Rousseau (LAL-Orsay)**



David Rousseau, SW news , S&C Plenary,  
4 April 2011

# Recent releases



- ❑ 15.6.X.Y :
  - "May reprocessing release"
  - Very low traffic these days (MC generators only)
  - Can still be used for "tail" of 2010, r15 analyses
- ❑ 16.0.X.Y
  - "Fall reprocessing release"
  - Frozen Tier0 policy
  - Low tag traffic
  - 16.0.X.Y will be used at Tier0 from beginning of 2011
  - Recommended for all pp analyses
- ❑ 16.1.X :
  - clone of 16.0.X.Y for TDAQ (like 15.7.0 vs 15.6.X.Y)
  - "2011 TDAQ release", for TDAQ expert only
- ❑ 16.2.X
  - clone of 16.0.X.Y for Heavy Ion
  - 16.2.2 used for HI reconstruction
  - "mail agreement" by HI experts to use mainstream release for future HI studies (starting with 17.X.Y)
- ❑ 16.4.X "Upgrade" digitisation/reconstruction clone from 16.0.2. To be merged in 17.0.0
- ❑ 16.5.0 "dev" release

# Releases



<https://twiki.cern.ch/twiki/bin/viewauth/Atlas/SoftwareReleasePlans>

<https://twiki.cern.ch/twiki/bin/viewauth/Atlas/DeliverablesForRelease17>

- ❑ 16.6.X.Y : (last AtlasProduction 16.6.3.4)
  - 2011 Tier0 release, also for MC digi/reco (MC10a), up to the summer
  - Fast reprocessing of first 2011 data imminent
  - "Frozen Tier0 policy" enabled (only tags with zero impact on output except for identified problems)
  - To be used also for MC11 Simulation
  - **Recommended for analysis on 2011 data / MC**
    - (should not be used on 2010 data because AODFix not guaranteed to give the same result)
- ❑ 16.7.0
  - Sw deadline 7<sup>th</sup> March 2011. Deployed 16.7.0.1 in particular.
  - "Dev" technical release
  - Lcg60 in particular (lcg60b integration a bit chaotic, non backward compatible root changes requiring changes on atlas change, and a tricky root bug at the end)
  - 16.7.0.1 deployed last week-end
  - To be validated, but not for real production. (Validation lower priority than finalisation of 16.6.X.Y, however to be done well before 17.0.0 is closed, to save time on 17.0.0 validation).
- ❑ 17.0.0
  - Sw deadline 2<sup>nd</sup> May 2011
  - From then on ID and Calo reco to be mostly frozen (to allow tuning of combined reconstruction).
    - In particular pixel clusterisation rework, or
  - First round of validation
- ❑ 17.0.1
  - Sw deadline 30<sup>th</sup> May 2011
  - All other sw frozen
  - release to be fully validated beg of July
  - Bulk reprocessing to start early september
- ❑ Please remember these deadlines in meetings/discussion this week
- ❑ **New r17 deliverable meeting Tuesday 12<sup>th</sup> April 4PM (tentative)**

# R17 General sw goals



- ❑ Release 17 will be used for summer reprocessing and Tier0 reconstruction from August till end of the year
  - Algorithm/calibrations to be adapted to experience of spring data taking (in particular high pileup and luminosity)
  - CPU/Memory/Disk size to be reduced as much as possible
- ❑ Other sw goals:
- ❑ Already agreed that in release 17 onwards only maintain the capability to:
  - read raw data bytestream from fall 2008 (large cosmics data taking).
  - read geant4 hits from march 2010 geant4 production onwards (release 15.6.3.10 )
  - read RDO/ESD/AOD and other derived format from Christmas 2009 reprocessing onwards (release 15.5.4)
  - Support geometry ATLAS-GEO-00-00-00 and higher (and ATLAS-CTB for calorimeter and ID passive material)
  - =>cleanup obsolete EDM classes and geometry
- ❑ Software quality (see next slides)
- ❑ Python cleanup (see next slides)

# R17 Python configuration



- ❑ Many reco and analysis jobs are using RecExCommon, which makes cleanup difficult, but time has come to do the following
  - disable standalone flags like "EvtMax", but force usage of `jp.AthenaCommonFlags.EvtMax` which is now auto-imported in athena
  - remove `PoolRDOInput`, `PoolHitsInput` etc.... (keep only `jp.AthenaCommon.FilesInput`)
  - autoconfiguration by default
  - disable outputs by default (`cbnt`, `tag`)
  - re-enable error detection in production
  - retire `doCommissioning`
  - decommission `GlobalFlags` (only leave `globalflags`)
  - Investigate properties locking to avoid override of configurations between domains
  - Finalise "ConfigurableFactory" to ease share of tools
  - python cleanup (wim's athena -t tool being developed)
- ❑ More in reconstruction meeting tomorrow

# Software quality



- ❑ Software quality goals for r17
  - cleanup all compilation warnings
  - cleanup checkreq warning (including public vs private use)
  - cleanup coverity warning
  - python cleanup (locking plus wim's athena -t tool being developed)
  - Post r17:
    - turn compilation warning into error integrate checkreq in gmake
    - enable floating point exception (unless fpe WARNING can be systematically reported)
    - Explore <http://code.google.com/p/include-what-you-use/>
    - In general : move from cycle “developer makes a mistake, put tag in nightly, mistake detected in nightly, something/someone pings the developer, which fixes his mistake” to “developer makes a mistake, gmake fails, developers fixes mistake, put tag in nightly”
- ❑ Cleanup savannah bugs (more than 200 open bug reports in several projects)
- ❑ In practice how to do this? Would like to propose a “spring clean-up week” (better name welcome).
  - Week of 18<sup>th</sup> 22<sup>nd</sup> April
  - Focus on release 17
  - All developers/all coordinators focus on the clean-up. Some coordinator of large projects will need help.
  - All tags submitted to 17 be labelled “clean-up”

# 64 bit build



- ❑ 64 bit validation (of 16.0.3.4) on going with low priority but converging
- ❑ ...now moving to 16.6.3
- ❑ Can we consider moving completely to 64 bit build ?
- ❑ Many plusses:
  - Less build combination to handle (note that slc6 coming)
  - No need to switch to a different build for high memory job (in particular high pileup digi/rco)
  - 10-20% CPU gain (maybe more if we are focussing our optimisation on 64 bit build)
- ❑ One big minus
  - Memory use times 1.5, so reco memory will be ~3.5 GB

# Life with pileup



- ❑ Chamonix meeting early february: “fat” bunch mode of LHC confirmed
- ❑ Some small dedicated production of up to  $\mu = 15$  (average number of collision per beam crossing) RDO
- ❑ Report at TOB 21 February :
  - Up to 60 s per event reco, 600kB AOD, 3500kB ESD
  - Panic!, lot of work by many people...
- ❑ Report at Atlas Weekly 29<sup>th</sup> March: First 2011 run with  $\mu=8$  :
  - reco CPU  $\leq 15s$ , AOD size 180 kB, ESD size 1.1kB
  - (more in EMB/reco meeting)
  - not the end of the story of course...



# Release deletion



- ❑ New batch of release deletion essentially done (also obsolete DBRelease)
  - ❑ “Release” area at CERN more and more different from grid sites (some releases not distributed, some build not distributed, DBRelease not deleted)
- => Would need regular report from a “typical” grid site
- ❑ Still I’m told that sites Atlas SW area still very close to 300 GB
  - ❑ New cycle of release deletion, pending full deployment of CVMFS
  - ❑ See next slides. Note:
    - No general call for feedback yet
    - Need a strong push moving away from all 15.6.X except 15.6.14



<http://atlas-computing.web.cern.ch/atlas-computing/projects/releases/releaseTree.html>

Release	Size
16.7.0	9G
16.6.3	9G
16.6.2	15G
16.6.1	10G
16.6.0	12G
16.5.0	11G
16.4.1	10G
16.4.0	10G
16.3.0	11G
16.2.2	14G
16.2.1	13G
16.2.0	12G
16.1.2	21G
16.1.1	11G
16.1.0	11G
16.0.3	47G
16.0.2	18G
15.6.14	11G
15.6.13	18G
15.6.12	16G
15.6.11	13G
15.6.10	15G
15.6.9	34G
15.6.3	25G
<b>Total</b>	<b>371G</b>

@CERN

@CERN

@CERN

@CERN

@CERN

@CERN

### Recent ATLAS Releases

Calendar View

Age in weeks 0 1 2 3 4 5 6 7 8 9 10



Release 16.7.0	AtlasProduction-16.7.0.1	
Release 16.6.3	AtlasProduction-16.6.3.4	
	AtlasProduction-16.6.3.3	
	AtlasProduction-16.6.3.2	
	AtlasProduction-16.6.3.1	
Release 16.6.2	AtlasProduction-16.6.2.31	
	AtlasProduction-16.6.2.6	HiggsAnalysis-16.6.2.6.2
		HiggsAnalysis-16.6.2.6.1
	AtlasProduction-16.6.2.5	
	AtlasProduction-16.6.2.4	BTagging-16.6.2.4.1
	AtlasProduction-16.6.2.3	HiggsAnalysis-16.6.2.3.1
	AtlasProduction-16.6.2.2	
	AtlasProduction-16.6.2.1	
Release 16.6.1	AtlasProduction-16.6.1.1	
Release 16.6.0	AtlasProduction-16.6.0.4	
	AtlasProduction-16.6.0.3	
	AtlasProduction-16.6.0.2	
	AtlasProduction-16.6.0.1	TAGProd-16.6.0.1.1
Release 16.5.0	AtlasProduction-16.5.0.4	
	AtlasProduction-16.5.0.3	
	AtlasProduction-16.5.0.2	
	AtlasProduction-16.5.0.1	
Release 16.4.1		
Release 16.4.0		
Release 16.3.0	AtlasProduction-16.3.0.1	
Release 16.2.2	AtlasProduction-16.2.2.4	
	AtlasProduction-16.2.2.3	
	AtlasProduction-16.2.2.2	
	AtlasProduction-16.2.2.1	
Release 16.2.1	AtlasProduction-16.2.1.5	
	AtlasProduction-16.2.1.4	
	AtlasProduction-16.2.1.3	
	AtlasProduction-16.2.1.2	
	AtlasProduction-16.2.1.1	
Release 16.2.0	AtlasProduction-16.2.0.6	
	AtlasProduction-16.2.0.5	
	AtlasProduction-16.2.0.4	
	AtlasProduction-16.2.0.3	
	AtlasProduction-16.2.0.2	
	AtlasProduction-16.2.0.1	
Release 16.1.2	AtlasP1HLT-16.1.2.1	AtlasCAFHLT-16.1.2.1.2
		AtlasCAFHLT-16.1.2.1.1
Release 16.1.1	AtlasP1HLT-16.1.1.8	

Rel	Production	kJob	Decision (?)	Feedback (cache listed are not to be deleted)
16.7.0	Grid		Keep	Dev release to be validated
16.6.3	Tier0/MC		Keep	Tier0+MC10
16.6.2	Tier0/MC10a pre		keep	
16.6.1	Tier0		Delete	Short lifetime
16.6.0	Tier0		Delete	Short lifetime
16.5.0	Dev release		Delete	"dev" release made obsolete by 16.7.0
16.4.1	Upgrade		Keep	Not used outside CERN
16.4.0	Upgrade		Delete	Not used outside CERN
16.2.2	HI Repro+ana		Keep	HI reprocessing and analysis
16.2.1	HI Tier0+ana		Delete	
16.1.2	TDAQ 2011		Delete	Not used outside CERN
16.1.1	TDAQ 2011		Delete	Not used outside CERN
16.0.3	Ana		Keep	All caches used in production
16.0.2	Fall repro		Keep, del caches	16.0.2.3 4 6 used in production, 16.0.2.2 5 by users
15.6.14	Last 15.6.X rel		Keep	All caches used in prod
15.6.13	pp Tier0 +ana		Keep	All caches used in prod
15.6.12	MC10 Geant4		Keep	All caches used in prod
15.6.11	evgen		Keep	All caches used in prod
15.6.10	MC09 evgen		Keep, del caches	15.6.10.6 7 8 used in prod
15.6.9	(may repro)		Keep, Del caches	15.6.9.8 used in prod
15.6.3	MC09 Geant4		Keep, Del caches	15.6.3.10 12 13 14 17 18 used in prod

To  
be  
revised

# Software shifts



- ❑ A number of Cat 2 SW shifts are available (and lots of slots free in the coming week):
  - Reco and sw validation shift
  - Simulation shift
  - RTT
- ❑ Output of these shifts useful for all developers:
  - Savannah bug reports
  - elog: <https://prod-grid-logger.cern.ch/elog/ATLAS+T0+SW+Validation+Logbook/>
  - Mailing lists about nightlies status (any one can register or contribute) [atlas-sw-dev-nightly-validation@cern.ch](mailto:atlas-sw-dev-nightly-validation@cern.ch), [atlas-sw-prodcache-nightly-validation@cern.ch](mailto:atlas-sw-prodcache-nightly-validation@cern.ch)
    - For example “DEV VAL 3 unusable unless one uses SomePackage-01-02-03”

# Movements



- ❑ New simulation convenors: Philip Clarke and Daniel Froidevaux (a big thank to Adele Rimoldi and Charlie Young). (Simulation workshop Wed AM and Fri PM)
- ❑ Event Management Board : Vivek Jain joined RD Schaffer
- ❑ Search for r17 reprocessing sw release coordinator close to completion
- ❑ Software Infrastructure Team :
  - Krzysztof Ciba left end December
  - Samir Ferrag joined SIT (@50%) leaving end of December
  - Oleg Zenin for three months
  - 4 new shifters for release build
  - Still short of people for more exploratory tasks (more at SIT meeting)

# Other reinforcement needed



- ❑ Core job transform:
  - Specific transforms Simu\_trf, Digi\_trf, Reco\_trf well maintained
  - ...but core trf is not at all
  - ~30% involvement (need good python experience)
- ❑ Reconstruction integration:
  - Integrating reconstruction development of ~200 developers
  - Fire fighting mode since last summer:
    - pp data taking
    - Heavy ion preparation and data taking
    - Fall 2010 Reprocessing
    - AODFix
    - "Life without ESD"
    - "Life with pileup"
    - ...
  - Sw shifts functioning=>many bug reports
  - =>Need 2-3 people 30-50% FTE:
    - Maintenance of integration tests closest to what will be run
    - Python reco configuration
    - First line bug investigation
- ❑ Physics Analysis Tools, Event Management Board, Performance Management Board also very short

# Conclusion



- ❑ Release 17 first deadline 2<sup>nd</sup> May
- ❑ Please have this in mind in discussions this week, and report on Rel 17 deliverable meeting 12<sup>th</sup> April
- ❑ Reinforcement needed in several areas (even if you're not interested maybe some lab colleague would be ?)