



# ATLAS Computing Status

Eric Lançon  
Simone Campana  
Richard Mount





# Major Developments in LS1



# Major Developments During LS1 (Software)

## Reconstruction:

- Goal: speed up\* by a factor 2 to 3\*\*

**Goal exceeded**

## Full Simulation:

- Goal: speed up significantly
- But also improve fidelity

**Factor ~1.5 speed**

**Factor ~0.8 speed**

## Fast Simulation:

- Goal: new FastCaloSim “good enough for taus”

**Ready late 2015**

## Analysis Workflow (more centrally managed):

- New “xAOD”, ROOT and Athena readable
- New “train-based” Derivation Framework
- New Analysis Framework

**In use**

**In use**

**Use starting**

\* for a 14 TeV ttbar event with pileup of 40 and Run 1 geometry

\*\* reconstruction and simulation speedups were assumed in Run 2 resource planning

# Major Developments During LS1 (Distributed Computing)



SLAC

## Distributed Data Management:

- DQ2 → Rucio (scalability and function for Run2 and beyond)

**In operation 1<sup>st</sup> Dec**

## Distributed Workload Management:

- ProdSys → ProdSys2
  - JEDI (Job Execution and Definition Interface)
  - DEFT (Database Engine For Tasks)

**In operation 1<sup>st</sup> Dec**

## New Data Management Strategy:

- All data assigned a type-dependent Lifetime
- Disk vs Tape residency decided by operations
- Expanded use of Tape (with Lifetime-based cleanup)

**Now being implemented (manual)**  
**Automated implementation in 2015**



# **Resource Usage in 2014**

**(Largely from the report to the RRB/C-RSG submitted mid February)**



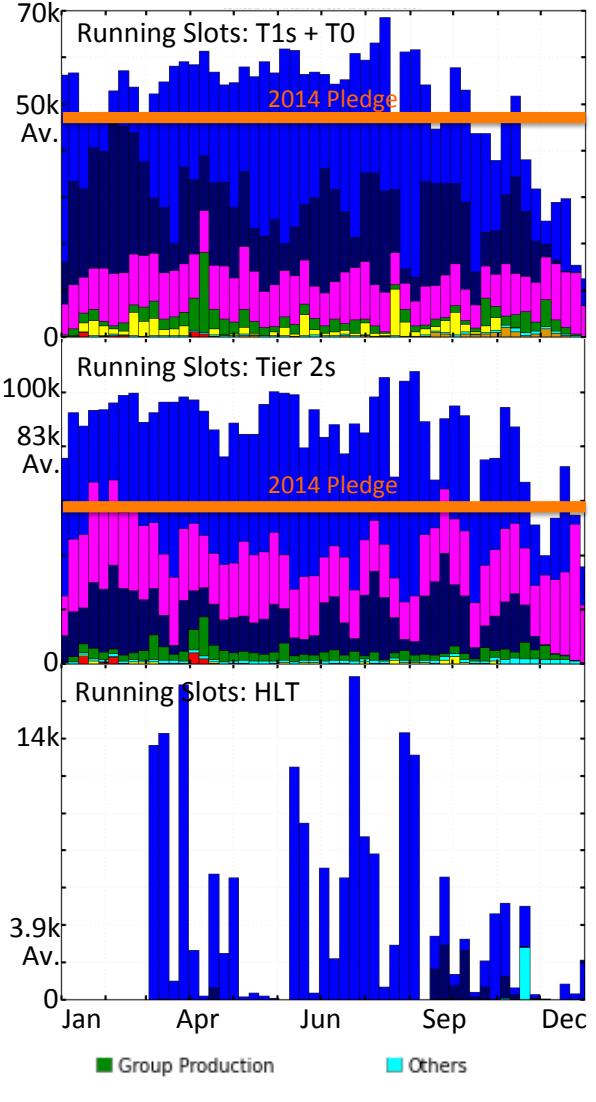

# Resource Usage in 2014

	<b>Pledges</b>	<b>ATLAS Max Available</b>	<b>ATLAS Used</b>	<b>CPU time/ Wall time</b>
CERN CPU (kHS06)	111	111	69	89.3%
Tier 1 CPU (kHS06)	355	-	381	90.8%
Tier 2 CPU (kHS06)	390	-	581	90.3%
HLT CPU (kHS06)		-	31	95.3%
<b>TOTAL CPU (kHS06)</b>	<b>856</b>	<b>-</b>	<b>1062</b>	
Tier 0 + CAF Disk (PB)	11	10	6	
Tier 1 Disk (PB)	38	41	36	
Tier 2 Disk (PB)	52	59	44	
<b>TOTAL Disk (PB)</b>	<b>101</b>	<b>110</b>	<b>86</b>	
Tier 0 Tape (PB)	27	-	30	
Tier 1 Tape (PB)	53	-	43	
<b>Total Tape (PB)</b>	<b>80</b>	<b>-</b>	<b>73</b>	


# Running and Pending Jobs in 2014



SLAC



# Running and Pending Jobs in 2014

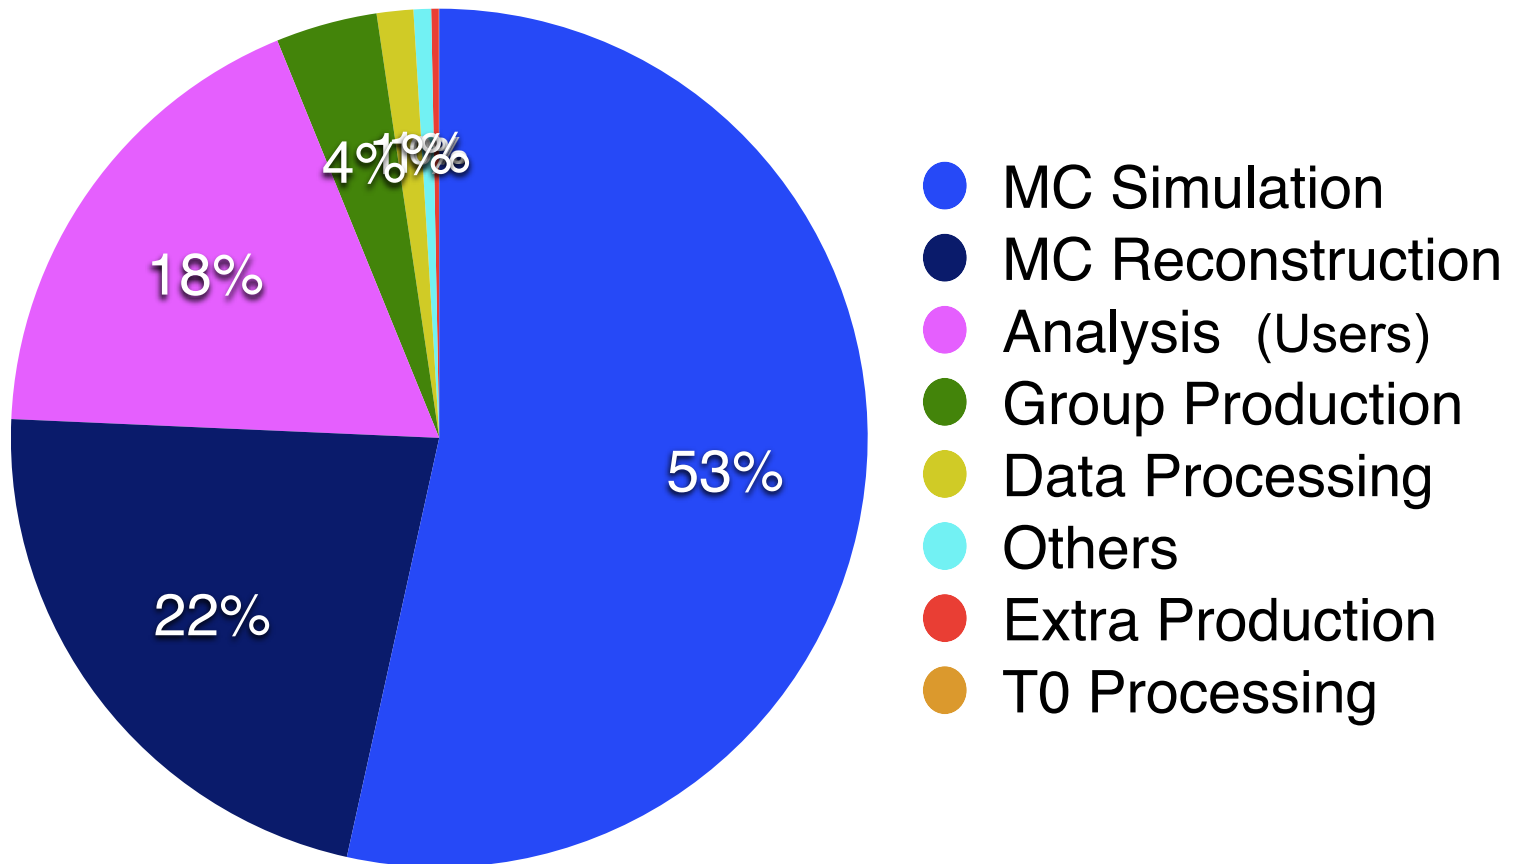


Temporary slowdown of production due to Rucio and ProdSys2 commissioning and new versions of ATLAS software


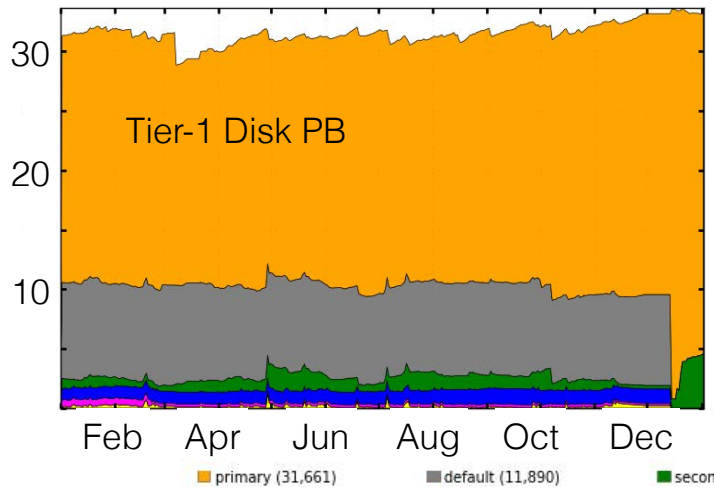




# Relative Wall Clock Time – T0, T1, T2

## Wall Clock consumption Good Jobs





# Disk Space Usage in 2014





# Disk Space Usage in 2014



SLAC




Rucio in use




# Disk Space Usage (end 2014)

Breakdown of overall disk usage at Tier0, Tier 1 and Tier 2




# Tape Usage in 2014



# Tape Usage Since January 1, 2015



Number of Physical Bytes (in TBs)  
54 Days from Week 00 of 2015 to Week 08 of 2015




Lifetime Policy

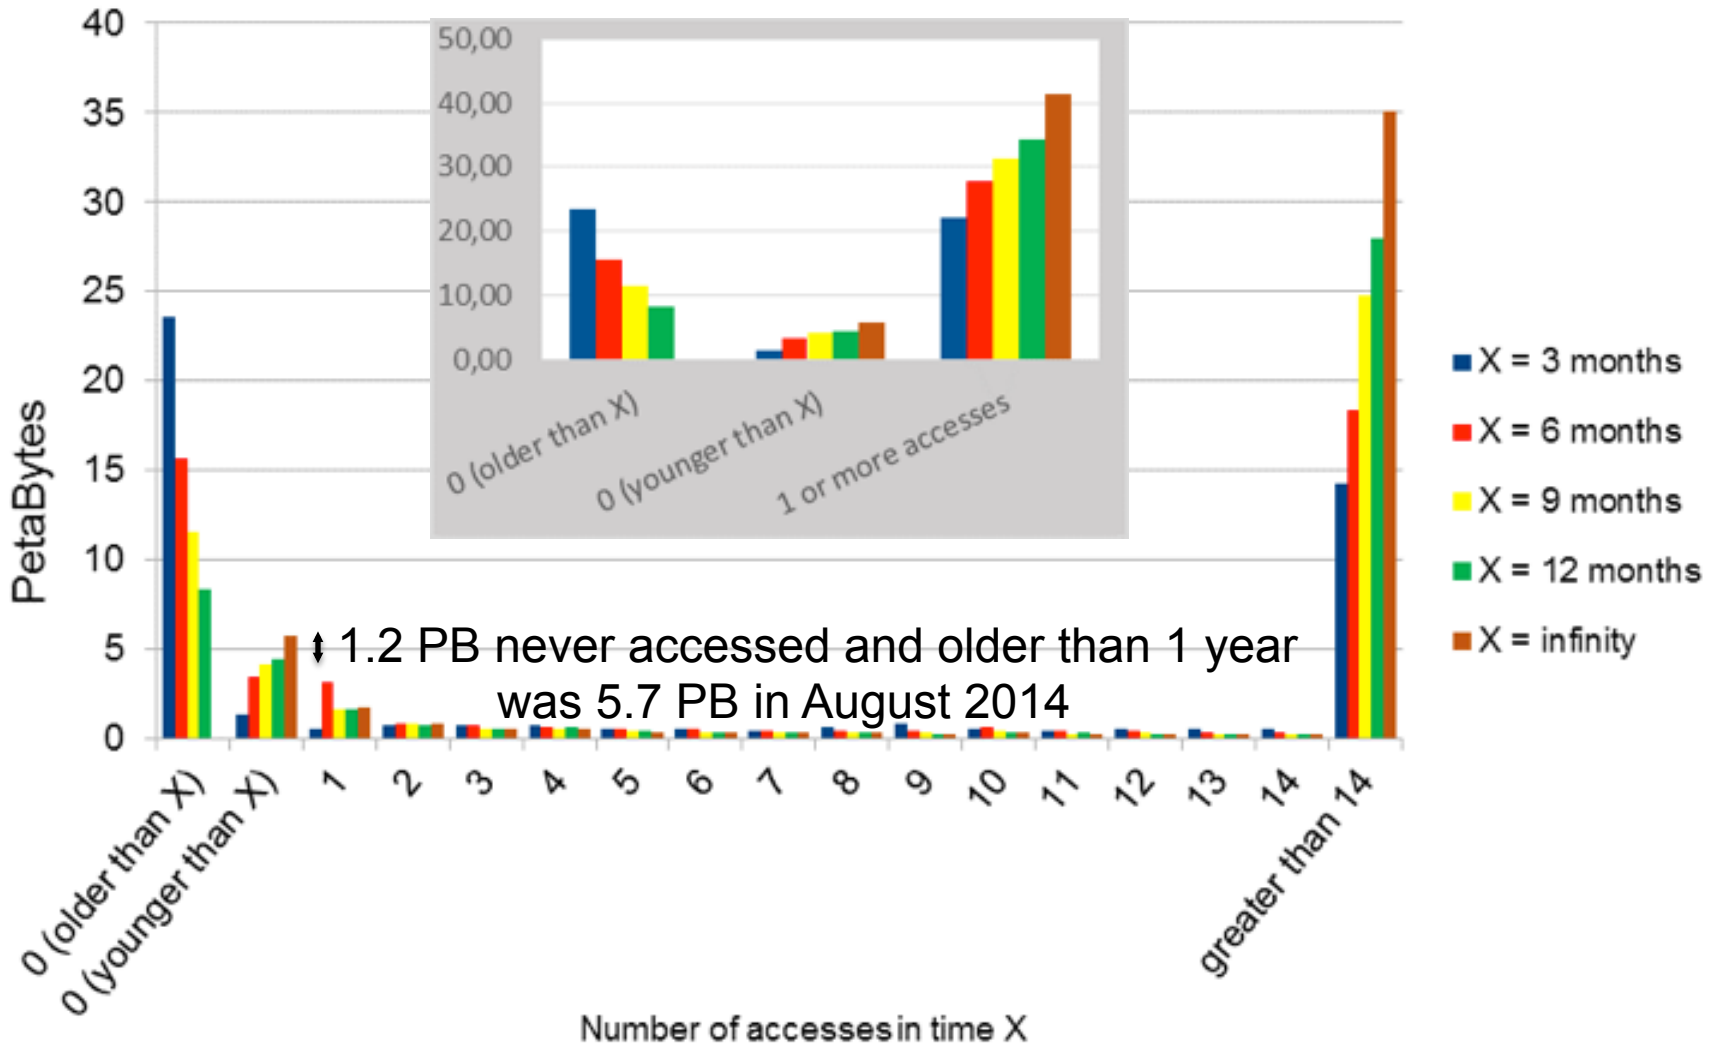
- |      |         |       |       |       |
|------|---------|-------|-------|-------|
| raw  | esd     | aod   | hits  | ntup  |
| desd | expired | desdm | user  | cbnt  |
| rdo  | d2aodm  | log   | draw  | group |
| dpd  | daod    | hist  | other | tag   |
| evnt | munt    |       |       |       |

Maximum: 73,363 , Minimum: 62,939 , Average: 70,858 , Current: 63,034

# Accesses to ATLAS DATADISK



# Accesses to ATLAS DATADISK







# **Resource Needs for Run 2**

**(Largely from the report to the RRB/C-RSG submitted mid February)**

# Significant Revisions (with respect to August 2014 submission)



SLAC

1. More realistic data “lifetimes”
  - Earlier planning had implicitly assumed very rapid physics analyses and very short lifetimes for the data used by the analyses. (e.g. no Run1 analysis after December 2015)
  - Reality, ATLAS published 106 papers in 2014 and more than 100 are in preparation now.
  - Can achieve adequate lifetimes at lowest possible cost by adding tape storage.
2. Steadily increasing the number of simulated events
  - Earlier planning had decreased the number of simulated events from 2014 to 2015 and then held it constant each year, to keep costs low
  - Revised planning (taking into account hardware costs and replacement profile) requests steady increase each year.
  - ➡ Greater CPU need provided entirely by beyond-pledge resources
  - ➡ Greater storage need provided entirely by tape

# Lifetime-based cleanup of Run-1 simulation

MC to be deleted	Lifetime [months]	T0 DATADISK [TB]	T1 DATADISK [TB]	T2 DATADISK [TB]	T0 TAPE [TB]	T1 TAPE [TB]	TOTAL DELETED [TB]
ESD	12	510	359	263	0	142	1,274
RDO	2	0	101	1,404	1	323	1,829
HITS	12	58	758	0	244	7,473	8,533
AOD	12	57	270	620	77	2,240	3,264
DAOD	12	0	1	0	0	0	6
Others	12	43	985	23	0	657	1,708
	Total	668	2,474	2,310	322	10,835	16614


# Lifetime-based cleanup of Run-1 Real Data

Data to be deleted	Lifetime [months]	T0 DATADISK [TB]	T1 DATADISK [TB]	T2 DATADISK [TB]	T0 TAPE [TB]	T1 TAPE [TB]	TOTAL DELETED [TB]
ESD	3	0	21	0	0	6	27
AOD	24	0	3	0	20	0	23
ESDAOD_hi	24	36	13	0	200	482	731
ESDAOD_sm	24	2	2	2	368	27	401
DESD,DAOD	24	5	227	5	530	113	880
NTUP	12	0	1	0	133	3	137
Others	various	9	27	0	300	1	337
Total		52	293	7	1551	632	2,536

# Original Run 2 Simulation Request



SLAC




**2014 Reality**


# Revised Run 2 Simulation Requirements



SLAC



# Revised Run 2 Simulation Requirements



Assumes FastSim good enough for most ATLAS analyses by 2016

# Run 2 CPU, Disk and Tape Requirements



SLAC

	2015 C-RSG	2016 C-RSG CERN-RRB- 2014-079	2016 ATLAS (Revised)	2017 ATLAS (August 2014)	2017 ATLAS (Revised)
T0 CPU	205	257	257	273	273
T1 CPU	450	540	540	691	691
T2 CPU	520	590	590	732	732
T0 Disk	14	17	17	19	19
T1 Disk	36	47	47	58	58
T2 Disk	53	72	72	98	98
T0 Tape	33	42	42	54	54
T1 Tape	65	84	116	108	209



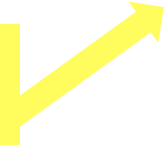
# Run 2 CPU, Disk and Tape Requirements



SLAC

	2015 C-RSG	2016 C-RSG CERN-RRB- 2014-079	2016 ATLAS (Revised)	2017 ATLAS (August 2014)	2017 ATLAS (Revised)
T0 CPU	205	257	257	273	273
T1 CPU	450	540	540	691	691
T2 CPU	520	590	590	732	732
T0 Disk	14	17	17	19	19
T1 Disk	36	47	47	58	58
T2 Disk	53	72	72	98	98
T0 Tape	33	42	42	54	54
T1 Tape	65	84	116	108	209

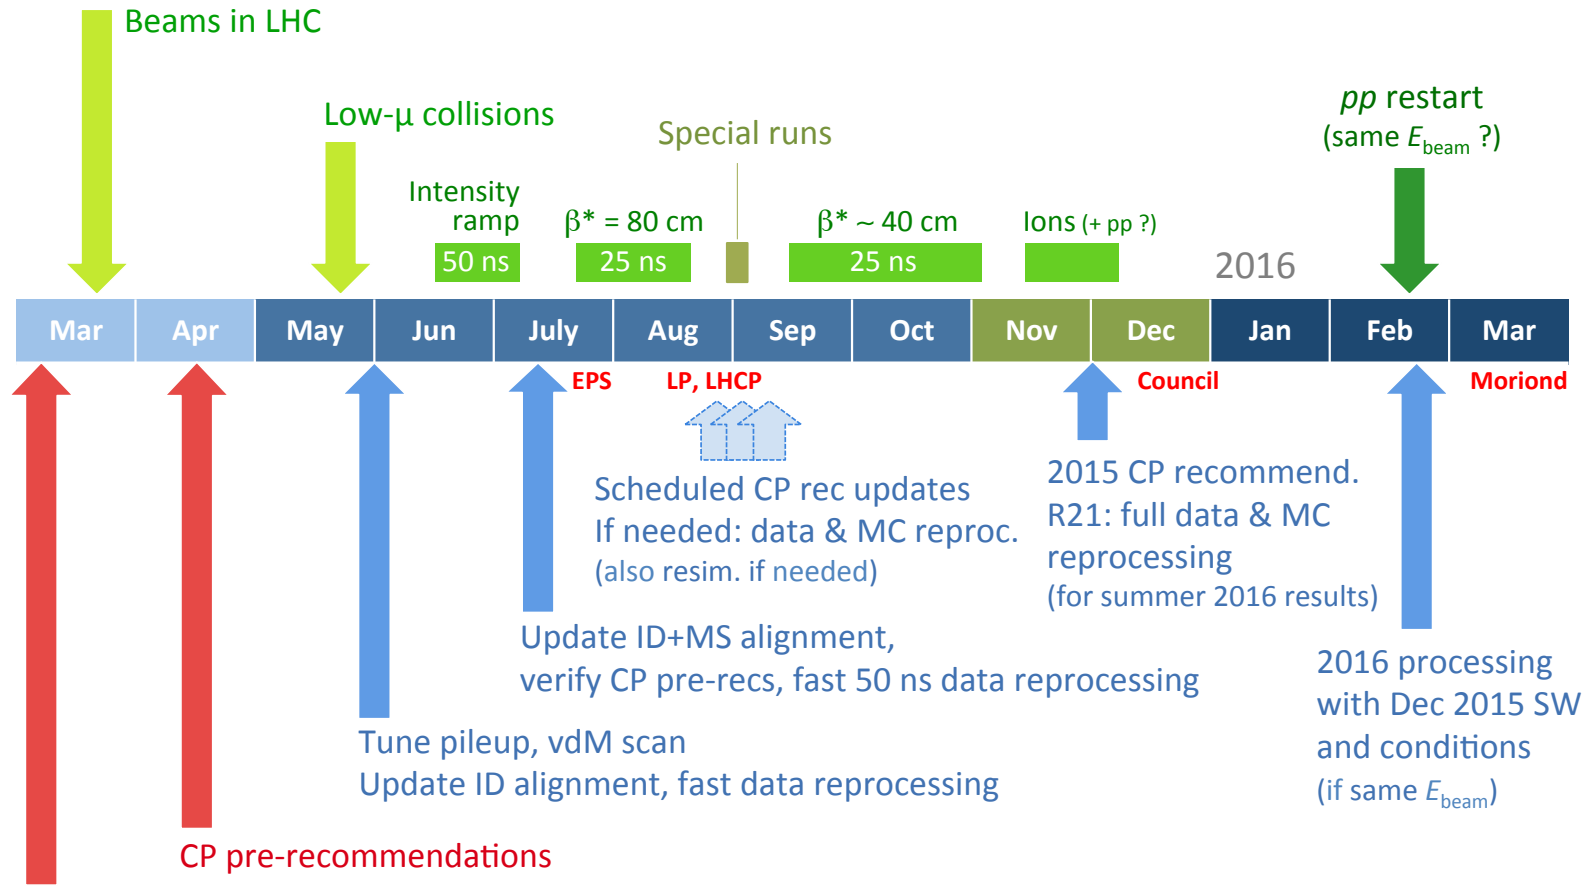
~half increase is for Lifetime  
~half increase is for Simulation



# ATLAS Goals for Unpledged Resources

	2015 Additional Goal	2016 Additional Goal	2017 Additional Goal
CPU [kHS06]	470	800	1100
Disk [PB]	0	0	0
Tape [PB]	at least 11	TBD	0

# Schedule – Coming 12 Months



R20.1 validated: launch MC15a digi+reco (2 × for 25 ns & 50 ns conditions)



## Inter-Experiment Development

- The present focus is ATLAS readiness for Run 2
- The Run-3-and-beyond focus is on efficient software development to support efficient and productive physics
- Skilled software development effort is a scarce and precious resource
- Common development wherever this can reduce the needed ATLAS effort and meet ATLAS needs

## ATLAS and the HSF

- ATLAS intends to take advantage of the HSF as a forum for identifying and fostering potential collaborative projects and presenting a coherent view of HEP computing to the wider world.
- ATLAS is particularly interested in collaborating on development of software in areas where it seems that through collaboration the products achieved would have better performance, be more effective in terms of cost and/or effort required and/or be better supported.
- ATLAS relies on existing common products such as G4, ROOT, LCG externals in general, generators and sees value in the HSF aiming at supporting those projects.

# Illustrations

ATLAS (or ATLAS members) effort contributing to software used beyond ATLAS:

- Geant4
- Sherpa generator
- TMVA
- RooFit, RooStat
- pyROOT
- Hammercloud
- PanDA
- Rucio
- ...

HSF Workshop January 2015

## Illustrations

Examples of ATLAS use of software developed by other experiments:

- Gaudi
- Gaudi HIVE (likely)

HSF Workshop January 2015

# Illustrations

Examples of ATLAS use of HEP-supported software:

- Generators
- ROOT
- Geant4
- CVMFS
- (HEP-supported) Cmake
- Grid Middleware

HSF Workshop January 2015



# Computing for Run 3 and beyond

1. Support for Run 3, 4 detector and physics studies
  - Simulation-focused “Software for Upgrade” group is operational
  
2. Development of Run 3,4 software (multithreading ...)
  - Studies ongoing (e.g. Future Framework Requirements)
  - Plans with milestones and resource needs later this year
  
3. Run 3, 4 computing resource needs model
  - In progress
  - Will be extension/extrapolation of the Run 2 model including data lifetimes

# ATLAS Summary

Developments in LS1 all aimed at maximizing physics productivity with achievable resource levels:

- Reconstruction speedup by more than a factor 3
- Full simulation speedup; ongoing fast simulation improvements
- New distributed computing software (more scalable, more efficient use of resources)
- New distributed data management strategy:
  - use tape storage (1/6 cost of disk) wherever it makes sense
  - Lifetime concept to remove least valuable data from storage

Revised resource request (more simulation, realistic data lifetimes):

- More tape (especially by 2017)
- More CPU → recognize the value of existing and future unpledged CPU

Common or collaborative software development

- Wherever it meets the needs at lower cost to ATLAS effort
- ATLAS supports the HEP Software Foundation

Planning for Run 3 and beyond

- Urgent and underway