

ATLAS Computing update

WLCG referees meeting

26 Feb 2019

James Catmore, Davide Costanzo

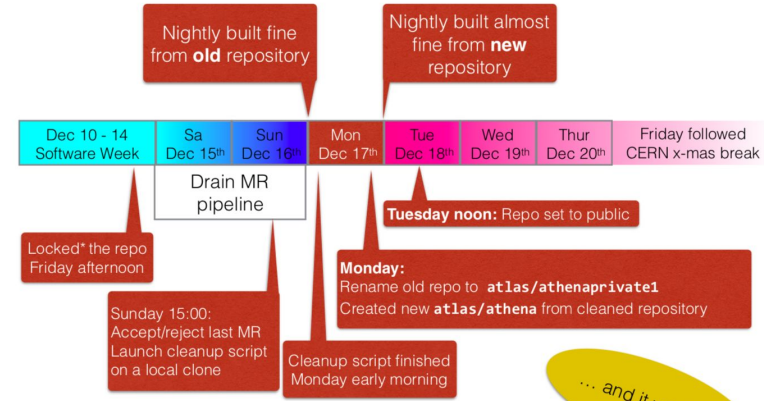
- Summary of the C-RSG document
 - report on 2018 usage of resources
- Work planned for LS2
- Initial comments on Run-3

Open software

- ATLAS software is now in a public repository
- © CERN for the benefit of the ATLAS Collaboration

The screenshot shows the GitHub interface for the ATLAS repository. At the top, a red banner indicates a GitLab outage: "GitLab unavailable for 5' around 19:00 tonight: <http://cern.ch/go/Hwz6>". The repository path is "atlas > athena > Repository". The current branch is "master" and the file is "LICENSE". A commit by Edward Moise is shown, dated 1 month ago, with commit hash 3d7e0604. The file size is 109 Bytes. The content of the LICENSE file is displayed as: "The software in this repository is released under the Apache 2.0 license, except where other licenses apply."



- Facilitate collaboration with computing scientists
- Developers can cite their code in the work and CV
- Executed after Run-2 and before the Christmas break to minimise disruption
 - General housekeeping and tidying up of repository



... and it worked!

2018 data taking

- Fantastic year of operation for LHC an ATLAS
 - 60 1 fb⁻¹ of “Good for Physics” data
- 103 paper submitted for publication on 2018
- 140 fb⁻¹ of data to analyse
 - Will drive computing usage during LS2



- Trigger rate vs month
- BLS menu extended in mid-July for R(K^{*}) measurement


Monte carlo simulation

MC16	Total number of events processed in 2018 (billions)
Event Generation	16.0
Simulation (FullSim)	9.9
Simulation (FastSim)	7.6
DigiReco	21.7

Data collected

Project	Stream	Events [M]	Volume [TB]
data18_13TeV (pp)	calibration	81752	2418
	physics	9832	9744
	other	121	134
data18_hi (Pb-Pb)	calibration	7863	2312
	physics	1245	1732
	other	24	41
data18_900GeV (pp)	calibration	5821	65
	physics	276	51
	other	0	0
data18_cos	all streams	356	167
other	all streams	50516	375

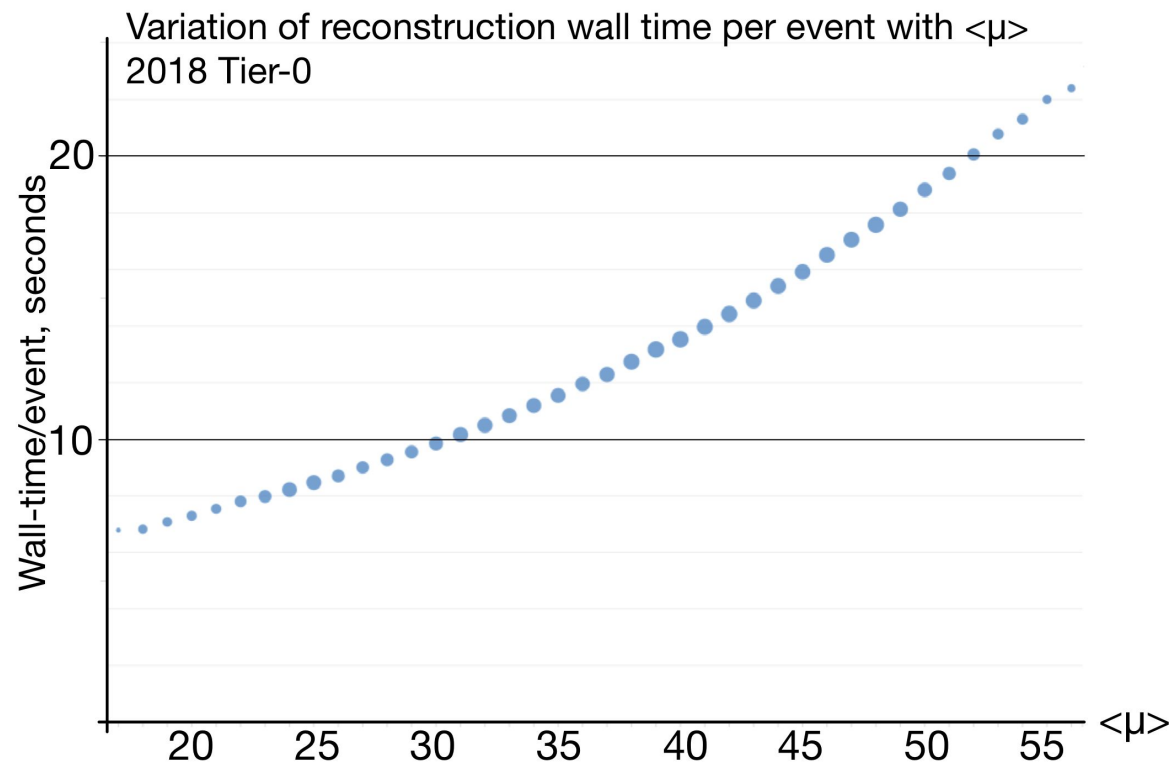
Summary of Tier0 operation in 2018



Heavy Ion.
Partial event stream

20% increase in May to
cope in case the 8b4e
scheme was used

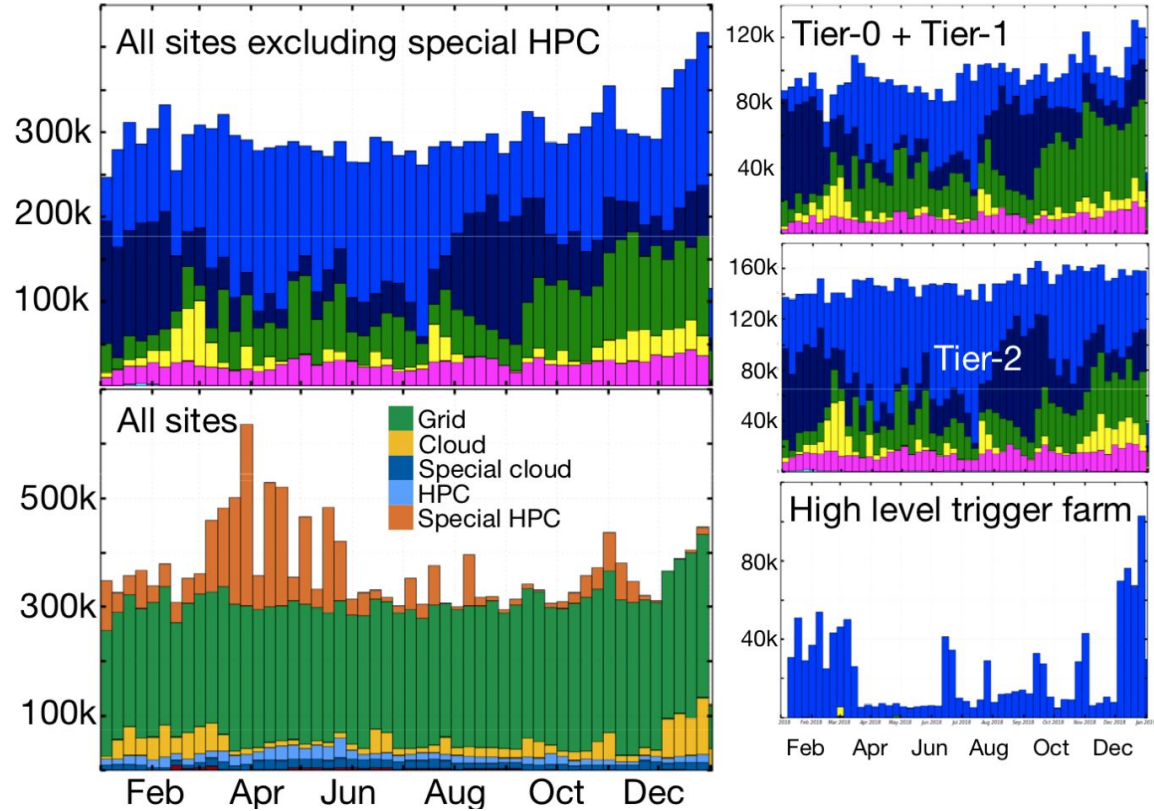
Tier0 always fully used



- $\langle\mu\rangle$ main parameter to establish Tier0 size
- Tier0 operating in single-thread mode
- Extrapolation to Run-3 points to a factor x2 increase
 - BUT reconstruction will be retuned for higher $\langle\mu\rangle$
 - We expect to need a factor x1.5 to x1.7 increase
- Spill over of prompt processing to grid used during HI run.

Jobs running on the grid in 2018

Running job slots




HLT farm used when not taking data. In particular a boost during the XMas break

Now offline for infrastructure work at P1 (no cooling)


HPC allocations at Cori and Titan used in 2018 (Note these are “weaker” nodes)

- Tier0: 496 kHS06 used vs 410 kHS06 pledged
- Tier1: 950 kHS06 used vs 952 kHS06 pledged
- Tier2: 1667 kHS06 used vs 1105 kHS06 pledged
- HLT farm: 235 kHS06 used (mostly for simulation)

Tier 1 CPU Consumption (Wall clock time) - 2018




Tier 2 CPU Consumption (Wall clock time) - 2018



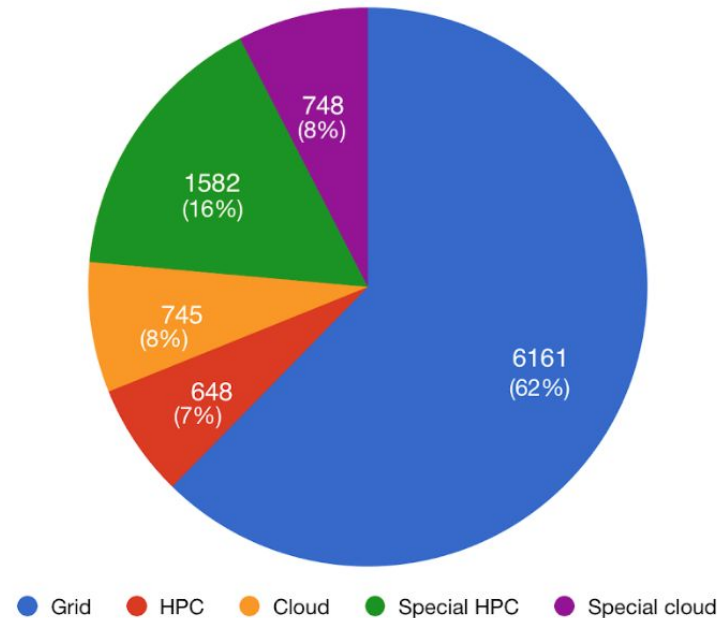
CPU usage distribution

- CPU usage dominated by MC Simulation, MC reconstruction and event generation
- Full Geant4 simulation runs on multiple resources
- Event generation time scrutinised for the MC workshop last November
 - several improvements identified


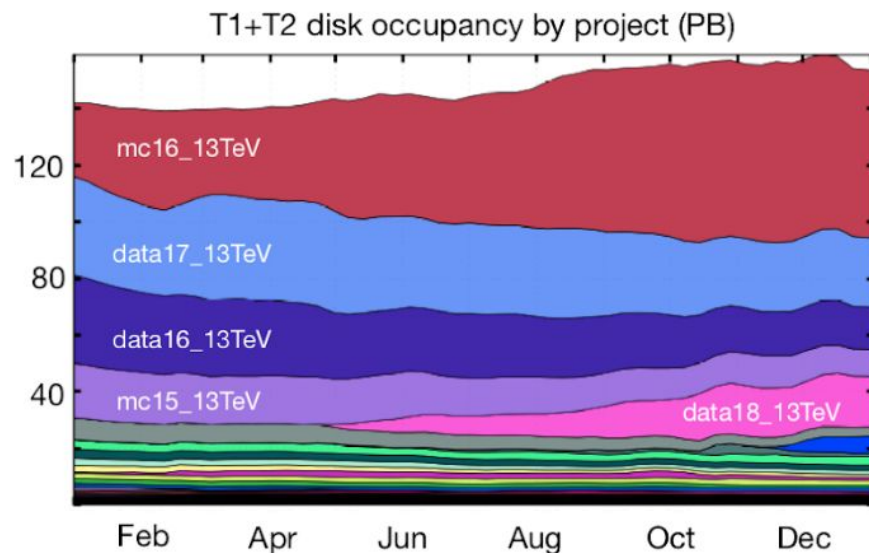
Wall clock consumption per workflow



Millions of events per resource (full simulation)

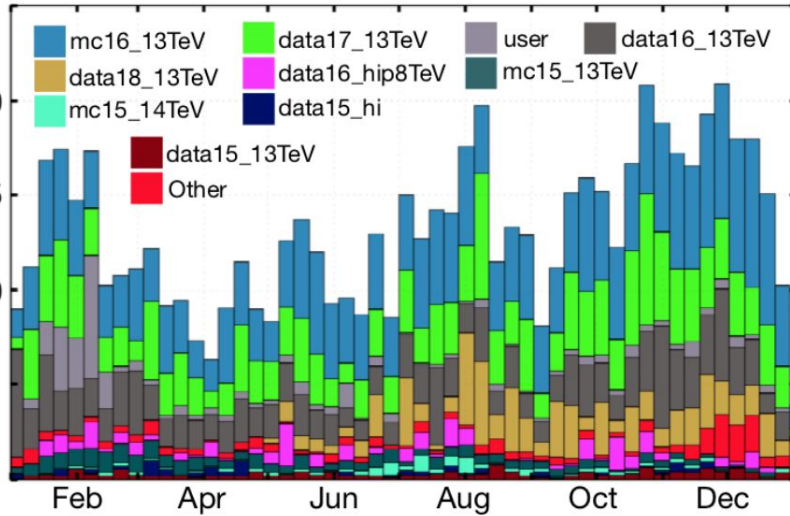


- Disks are fully occupied
- Space mostly used by analysis formats (AOD and DAOD) directly accessed by users
- Disk usage scrutinised by a Resource Management team
- Looking forward to 2019 pledges being delivered - mc16 running during LS2



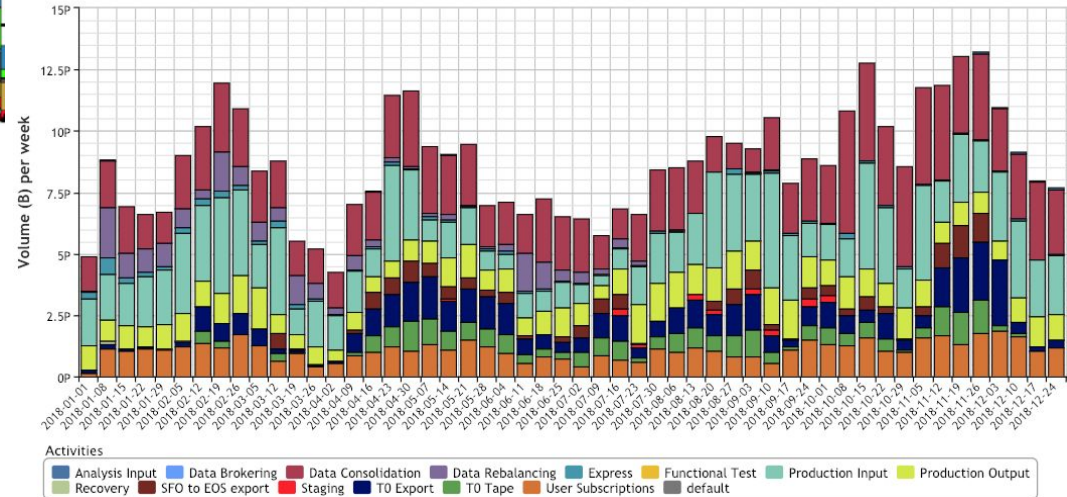
Volume of data processed and transferred

Data volumes processed per week by project, PB

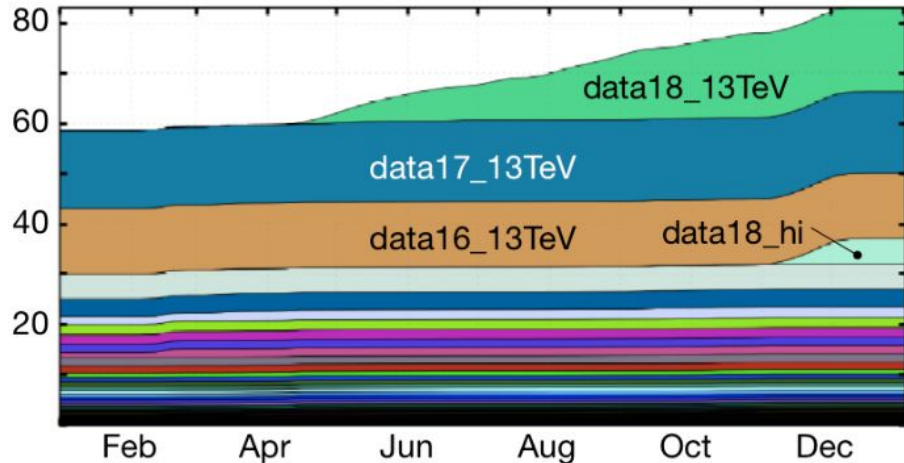


dashboard

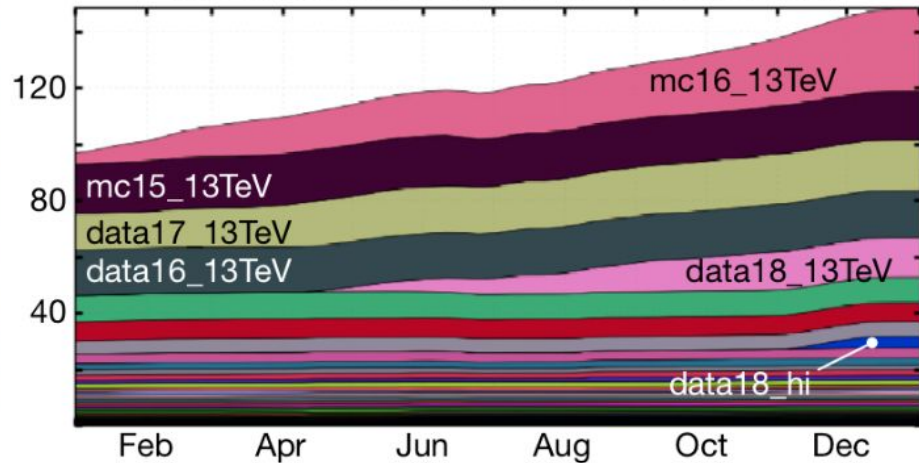
Transfer Volume
2018-01-01 00:00 to 2019-01-01 00:00 UTC



Tape occupancy at Tier-0 by project (PB)



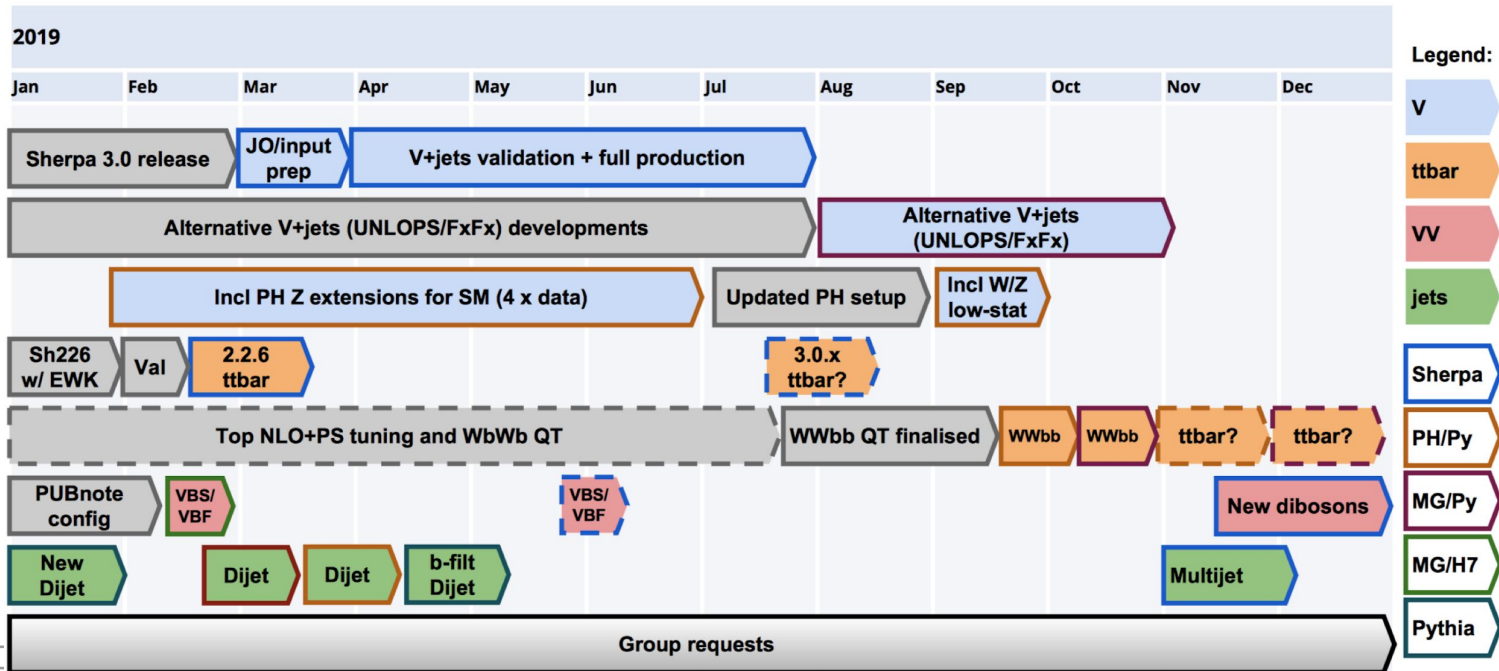
Tape occupancy at Tier-1 sites by project (PB)



Plans for LS2

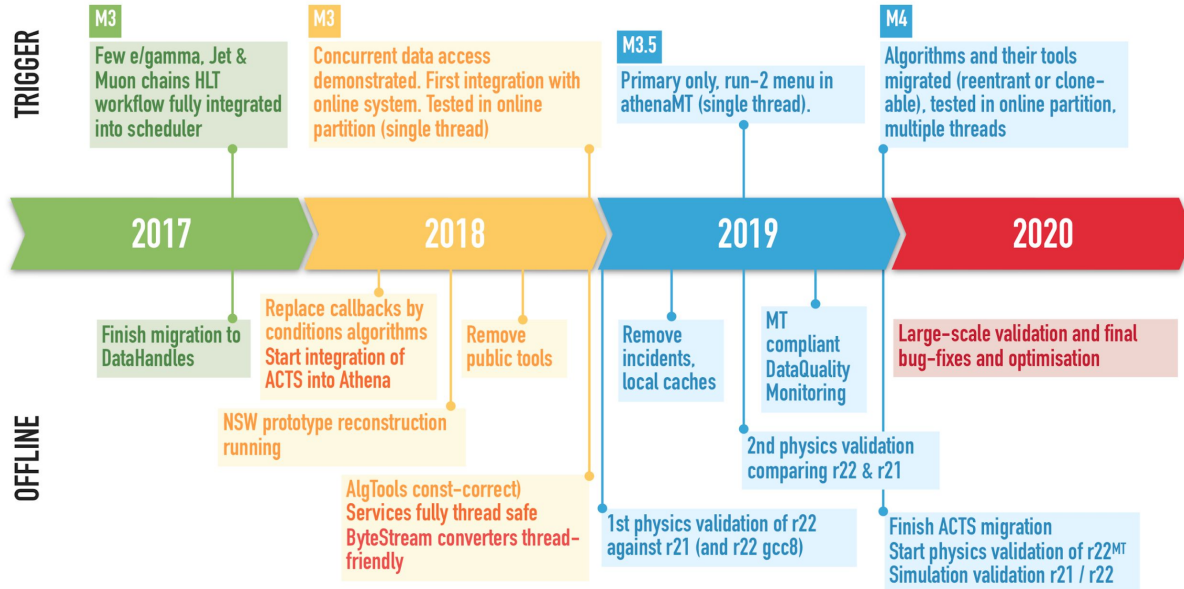
- Data reprocessing planned for specialised cases
 - eg BPhysics and Light State stream for a fraction of 2019, Heavy Ion
- Full reprocessing of 2017-18 data under discussion
 - No compelling reason at present. 2017 reprocessing and 2018 Tier0 reconstruction are good for physics

Monte Carlo simulation planning



Long Shutdown 2 (LS2) Software plans

- Updates software release plan.
 - Linked with Trigger milestones
 - More emphasis on validation
- Still understaffed. But we are seeing a reaction from the collaboration
 - Actively seeking new effort. A list of missing tasks was circulated to the collaboration



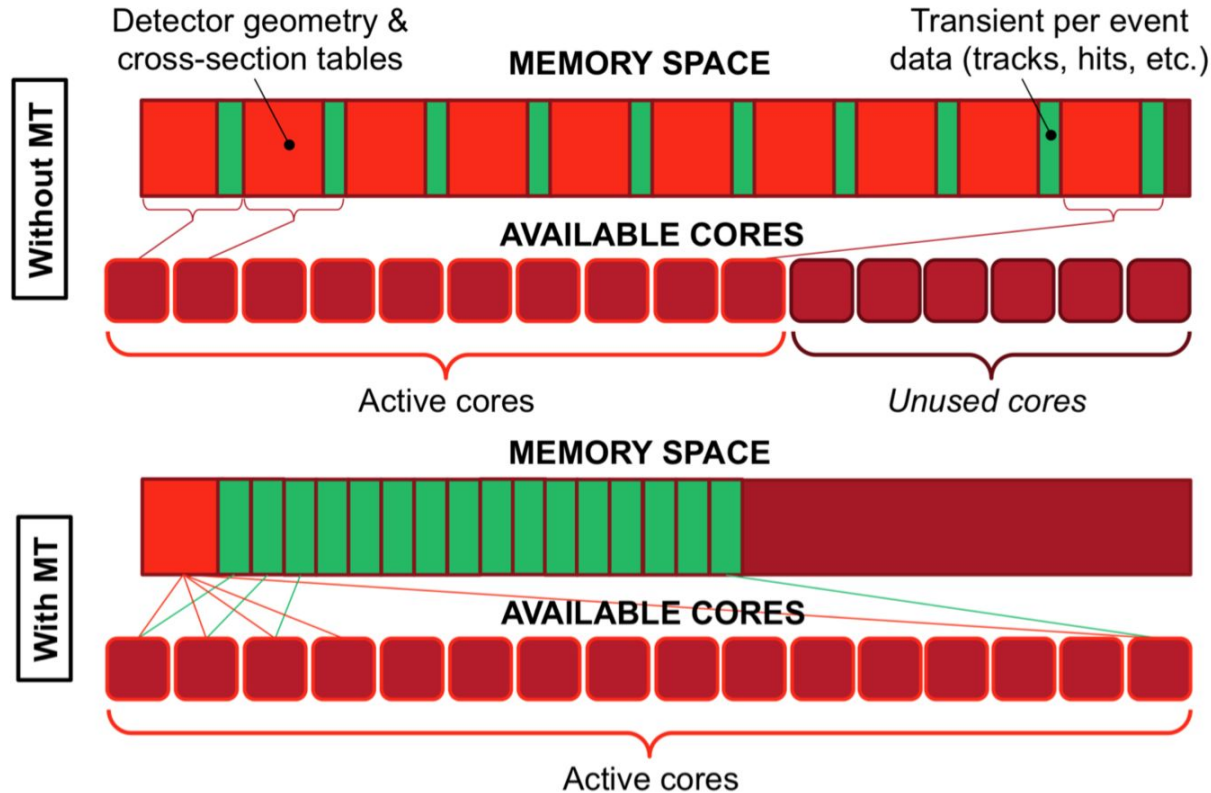
Analysis software deliverables	Run 3 physics	Run 4 development
AMSG 3 recommendations implemented leading to at least 30% disk savings	CRITICAL	CRITICAL
Review of EDM and possible simplifications	EXPEDIENT	EXPEDIENT

Simulation deliverables	Run 3 physics	Run 4 development
FastCaloSimV2 fully validated and ready for production	CRITICAL	CRITICAL
FastChain complete and ready for production	EXPEDIENT	CRITICAL
Multi-threaded simulation production-ready	EXPEDIENT	CRITICAL

Conditions / databases deliverables	Run 3 physics	Run 4 development
MT-compliant in-file metadata	EXPEDIENT	CRITICAL
PyCOOL and AMI fully compliant with Python 3	CRITICAL	CRITICAL
Oracle services compliant with new licencing regime	CRITICAL	CRITICAL
COOL-REST prototype ready for production	EXPEDIENT	CRITICAL

Reconstruction deliverables	Run 3 physics	Run 4 development
Release 22 fully validated and performant in single threaded mode (AthenaMP)	CRITICAL	CRITICAL
Integration of New Small Wheels	CRITICAL	CRITICAL
Integration of new tracking software (ACTS)	EXPEDIENT	CRITICAL
Release 22 able to run in multi-threaded mode (AthenaMT)	EXPEDIENT	EXPEDIENT

Common software deliverables	Run 3 physics	Run 4 development
Offline software fully Python 3 compliant	EXPEDIENT	CRITICAL
Configuration mechanism overhauled	CRITICAL	CRITICAL
New geometry framework ready for integration for Run-4	Not relevant	CRITICAL
AthenaMT performance optimised	EXPEDIENT	CRITICAL



- Simulation with Geant4MT operational
 - Technical issues being sorted
- Parts of the reconstruction workflows now work in MT
 - Goal to expand the test suite validation in H2 of 2019
- Tracking migration team
 - Evolving towards ACTS
 - Closely monitored

- CPU: We aim to produce 50% of our simulation using Fast Calo Sim v2
 - Faster simulation means faster turnaround time for MC samples
 - MC statistics was an issue at the beginning of Run-2
 - ATLAS will be more agile to changes in generators, simulation, etc
 - Overall more MC will be produced
- Disk: ongoing analysis model for Run-3 study group
 - Disk growth is a concern (as highlighted by the C-RSG)
 - We aim to reduce our disk footprint for Run-3
 - Preliminary recommendation to the collaboration in February. Final report in June
 - Move towards DAOD_PHYS and DAOD_PHYSLITE
- Tape: No issues in Run-2
 - About 75% of our tape storage used. Will catch up towards the end of 2020
 - Reprocessing of MC and data running from tape
 - We will request more tape in 2021
- Resources expectations
 - Below flat-budget in 2019, No increase in 2020
 - We expect to return to flat budget increases in Run-3
Eg (2021 resources) ~ 1.5x (2018 resources)
 - 2021 is a commissioning year, increase could be over 2021/22.
 - Still plan for a rich physics programme in 2021!