Software & Computing status

L. Poggioli, LAL Orsay

- Run-2 status
- Finalizing Run-2
- Preparing Run-3
- Towards HL-LHC

Based on
- Various S&C weeks
- ICB meetings
- March Sites Jamboree
- HSF/WLCG in Naples (see Catherine's talk)

Run-2

ATLAS pledges 2017

- 2017 (50% more data, 20% more CPU & Disk)

All ATLAS	Increase (2016->October requests)	Net increase for France	Fraction FR/all-ATLAS
T1 CPU KH	77%	36%	9.5%
T1 Disk pB	45%	13%	10.4%
T1 tape pB	62%	44%	9.5%
T2 CPU kH	99%	43%	7.4%
T2 Disk pB	15%	7%	9.1%

- Outcome (all ATLAS)
  - Shortage: CPU -14% Disk -2% Tape -5%.
  - #MC evts reduction, w/ 1kHz HLT & processing all events
- Ongoing optimization
  - Train production from tape,
  - AOD size reduction (~30%)
  - Workflow improvements

ATLAS pledges 2018

CERN-RRB-2017-057

Resource	Site	2017 Pledge	2018 ATLAS	Growth	2018 CRSG	Growth
CPU (kHS06)	T0+CAF	404	411	2%	411	2%
	T1	808	949	17%	949	17%
	T2	982	1160	18%	1160	18%
Disk (PB)	T0+CAF	25	26	4%	26	4%
	T1	69	72	4%	72	4%
	T2	78	88	13%	88	13%
Tape (PB)	T0+CAF	77	94	22%	94	22%
	T1	174	195	12%	195	12%

- ATLAS requests are within the expected flat budget increase and below the average 2013-2017 increase.
- CRSG recommends the requests.
- Beyond pledge resources is about 30% of the pledges, ATLAS expect to continue to receive a sizeable amount of over pledge CPU, which remain a risk for the experiment.

- Increase wrt 2017@T1 & T2s: in range 12-18%
- Except Disk@T1, T0: 4%

LHC/T0 planning 2018

LHC beam back 30/04/2018!!

- ATLAS baseline 2018 running scenario established, maximizing the physics within the various constraints
  - Based on LHC plan, 25ns BCMS, 2500b, 1.3e11 ppb, L=2.2e34
  - Will level lumi to 2.0e34 (mu = 56)
  - Trigger menu based on 2017 1.7e34 menu
  - Implications for Tier-0, Castor tape storage, data export, downstream resources, etc were part of the considerations
- Our Tier-0 processing model shows 20% capacity shortfall for processing all of physicsMain
- Grid spillover will be fully commissioned, validated and used in steady state during 2018 running for e.g. for B physics stream
- Spillover will require more operational effort, for both computing and data prep
  - Up to 0.5 FTE additional for ADC ops
  - Second reprocessing coordinator needed

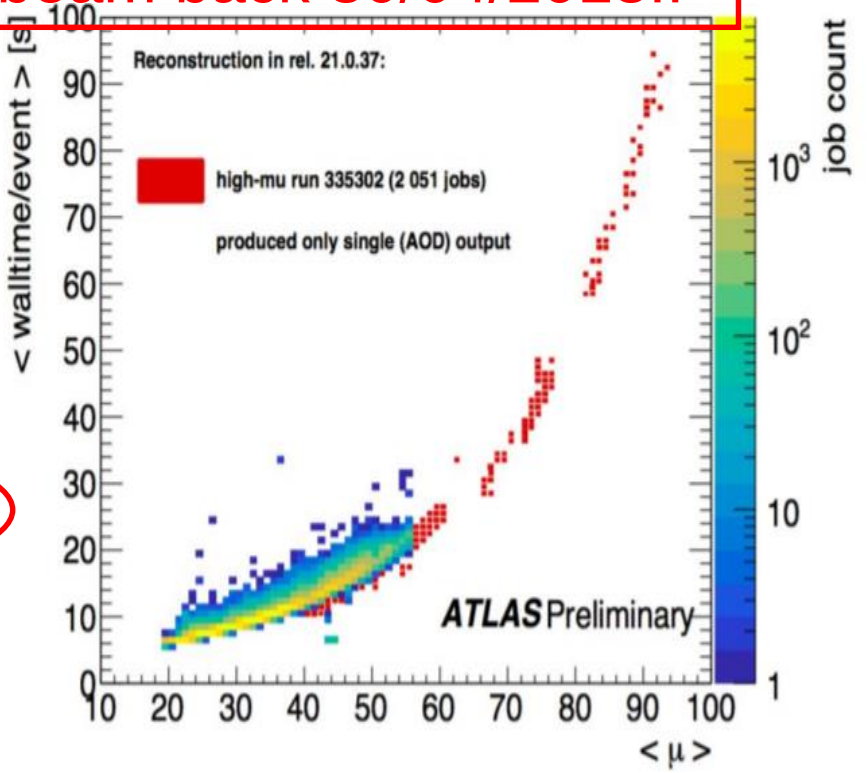

Strong wish to keep physicsMain processing at Tier-0 to ensure the timely integrity of the physics_Main data quality assessment cycle

CERN IT has responded favorably to our +20% Tier-0 CPU request

WORLD Cloud (1)

- Fully activated end March 2016
- Going definitely away from MONARC model
- Dynamic, tasks not confined to a cloud. Group of processing sites defined dynamically/task
- Task **nucleus**
  - Task brokerage choose nucleus for each task wrt data locality, queued work & available storage
  - T1s and bigger T2s are defined as nuclei
  - Output aggregated in task nucleus

WORLD Cloud (2)

- Task **satellites**
  - Run jobs and ship the output to the nucleus
  - Job brokerage selects satellites for each task, based on usual criteria (#jobs, data availability)
  - Satellites are selected worldwide: a network weight matches well connected nuclei & satellites
- Nuclei http://adc-ddm-mon.cern.ch/ddmusr01/NUCLEUS_DATADISK.html
  - Currently T1s and ~20% of T2s Better T2 disk usage!!
  - 65% datadisk in nuclei, aim to increase to ~80%
    - Today: CC, Tokyo, LAPP in FR-cloud

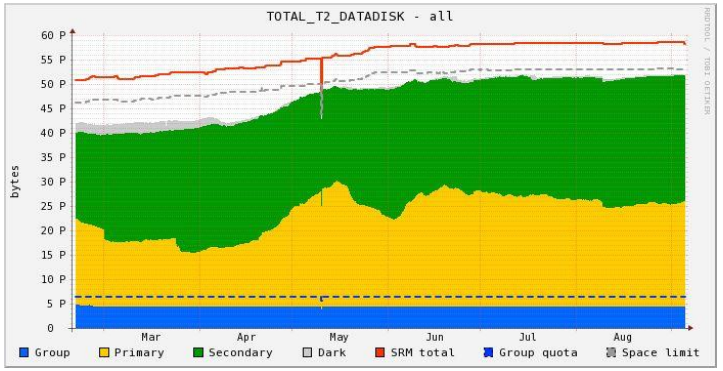

Resource usage in 2017 (1)

CPU availability in 2017

- Above pledges
- Dominated by Simulation
- MC limited stat is an issue (eg VH H to bb)

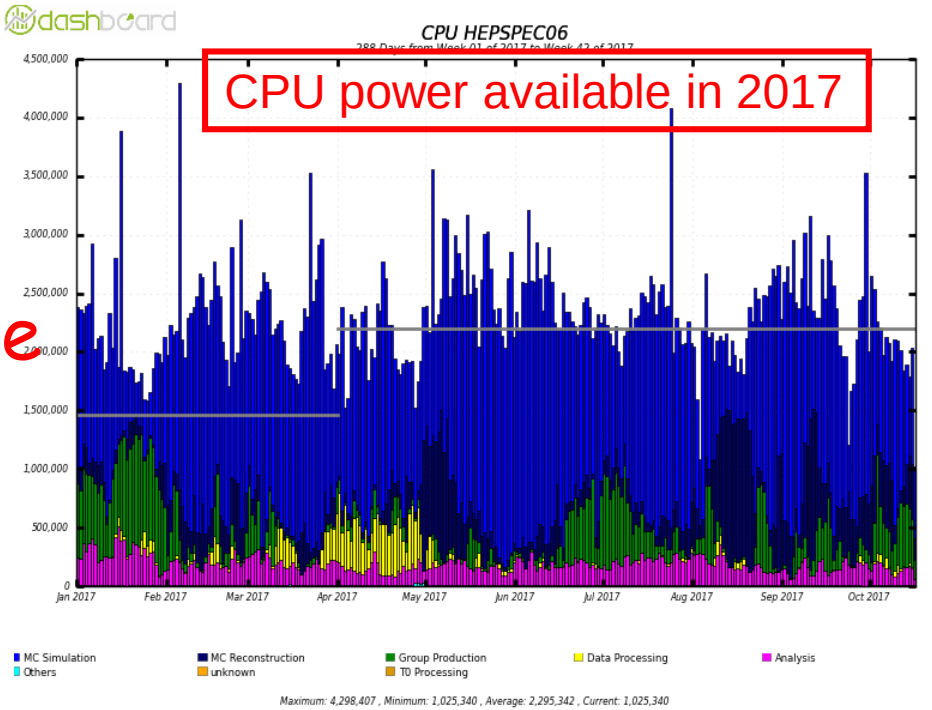

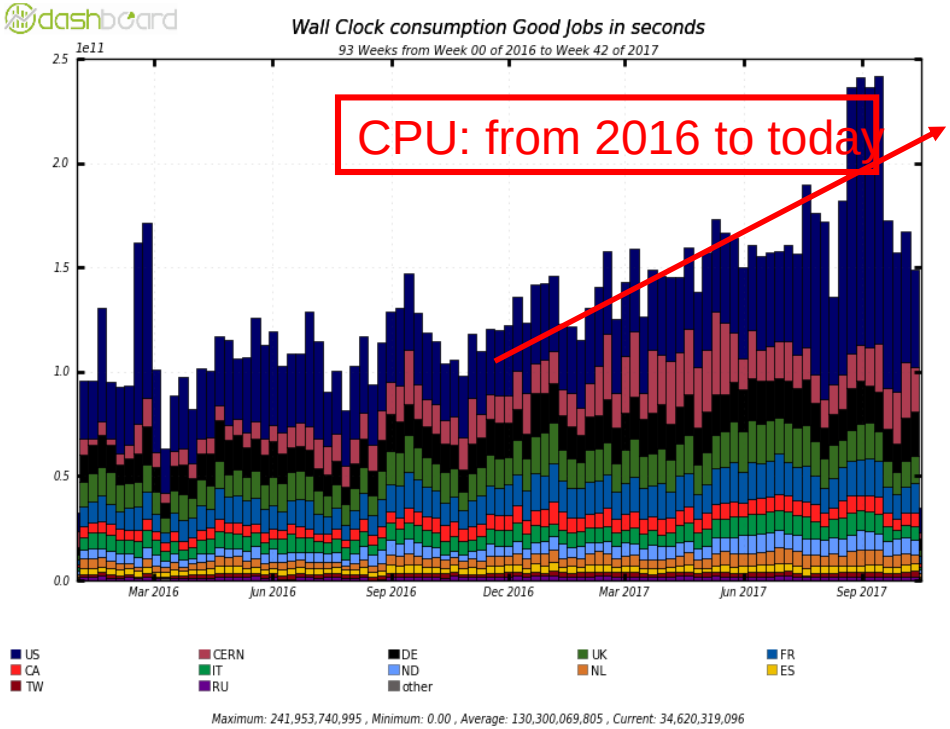

- 5-15 billion evts / day
  - few million/day simu.
- 2-5 PB of input/day
- >1 million jobs/day

Resource usage in 2017 (2)

Storage increase & usage

- T1s disks full at 85%
  - +5-10% for tape staging
- T2s disks full at 90%
  - Old problem solved
  (eg Thx Nucleus/Satellites Model)

eg CCIN2P3 DATADISK (2014-2017)

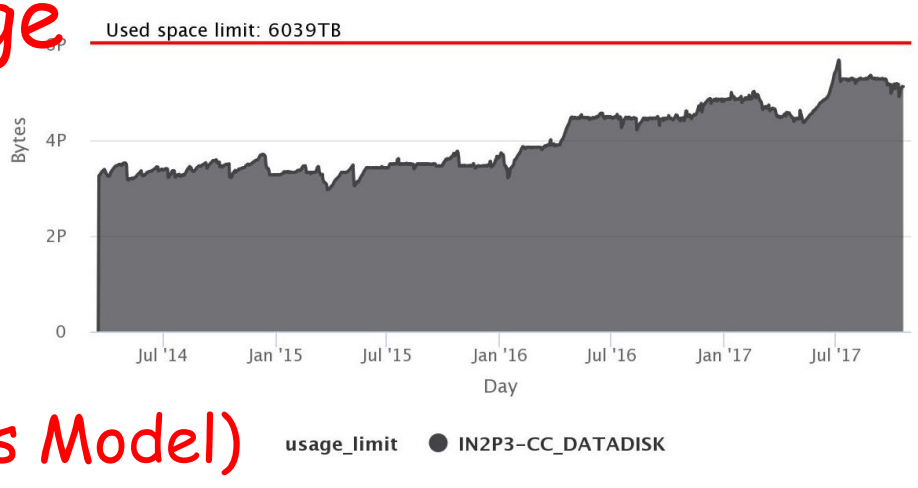

CCIN2P3 DATATAPE (2014-2017)

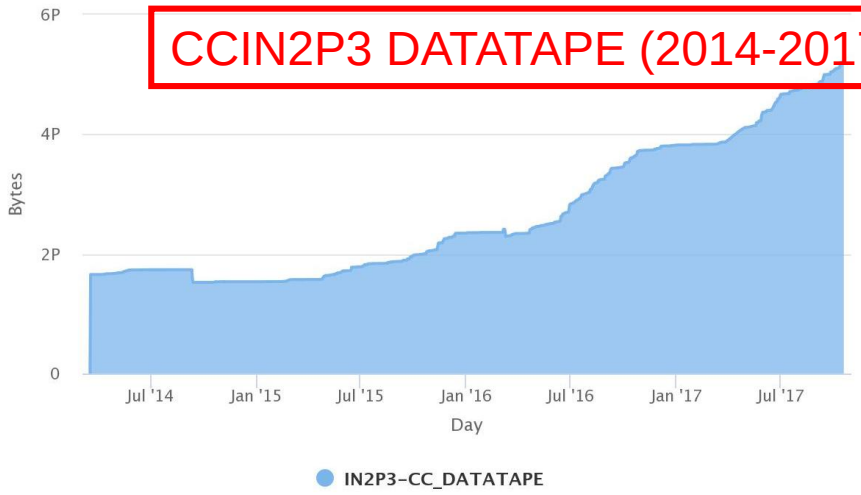

CCIN2P3 MCTAPE (2014-2017)

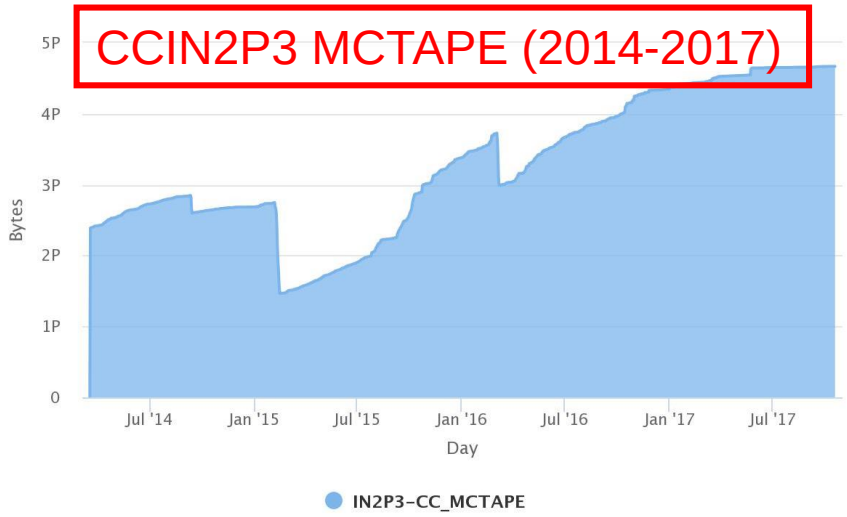

CPU Efficiency

Loss of CPU efficiency (CPU/WT) surveyed by LHCC

Introducing MCORE caused a loss of 10% efficiency. Some of this is real, due to serial operations in MCORE environment (e.g. I/O). Some is artificial (initialization accounted differently in MCORE and SCORE)

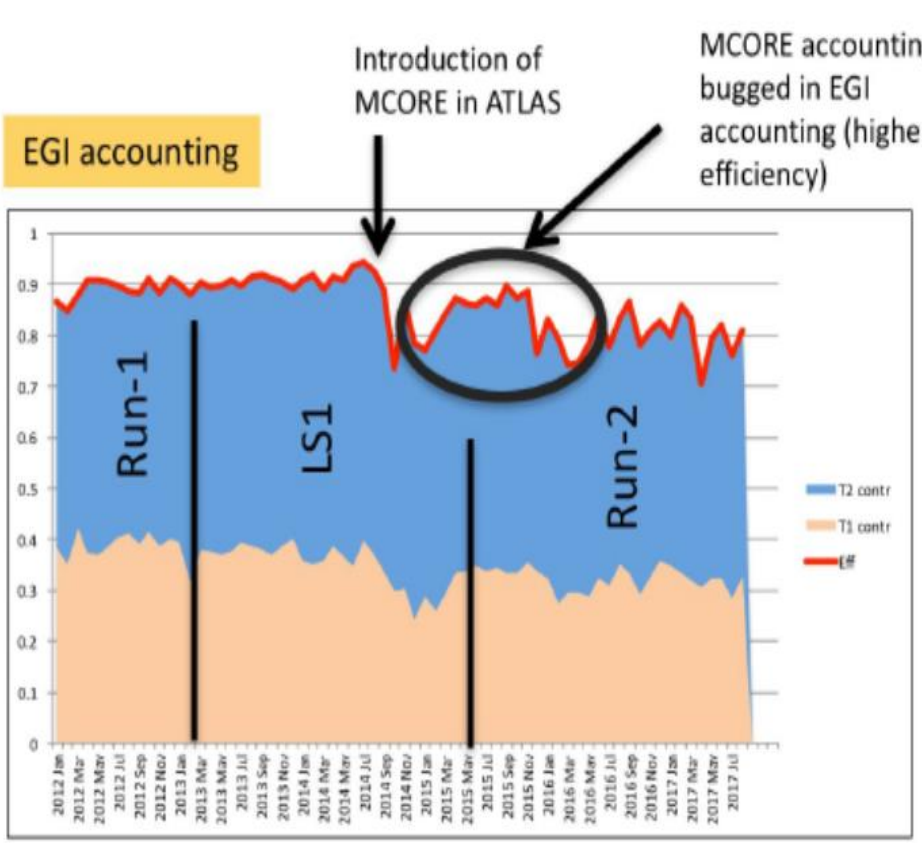

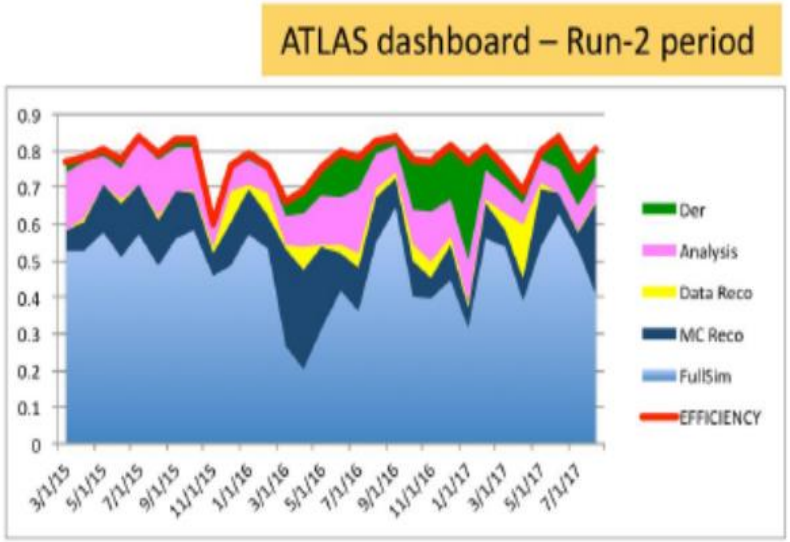

Conclusion-I: no obvious drop of ATLAS efficiency in Run 2

Conclusion-II: that does not mean we should accept 80% is good. We continue our work understanding performance and improving it

e.g. using checkpointed images to reduce serial init

Extra resource (1)

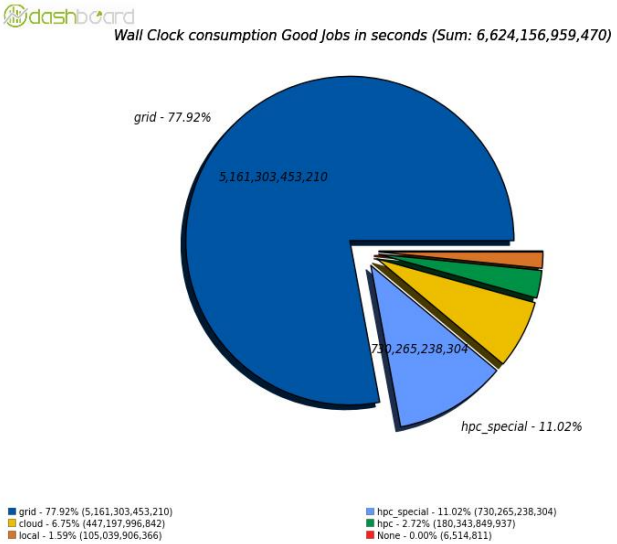

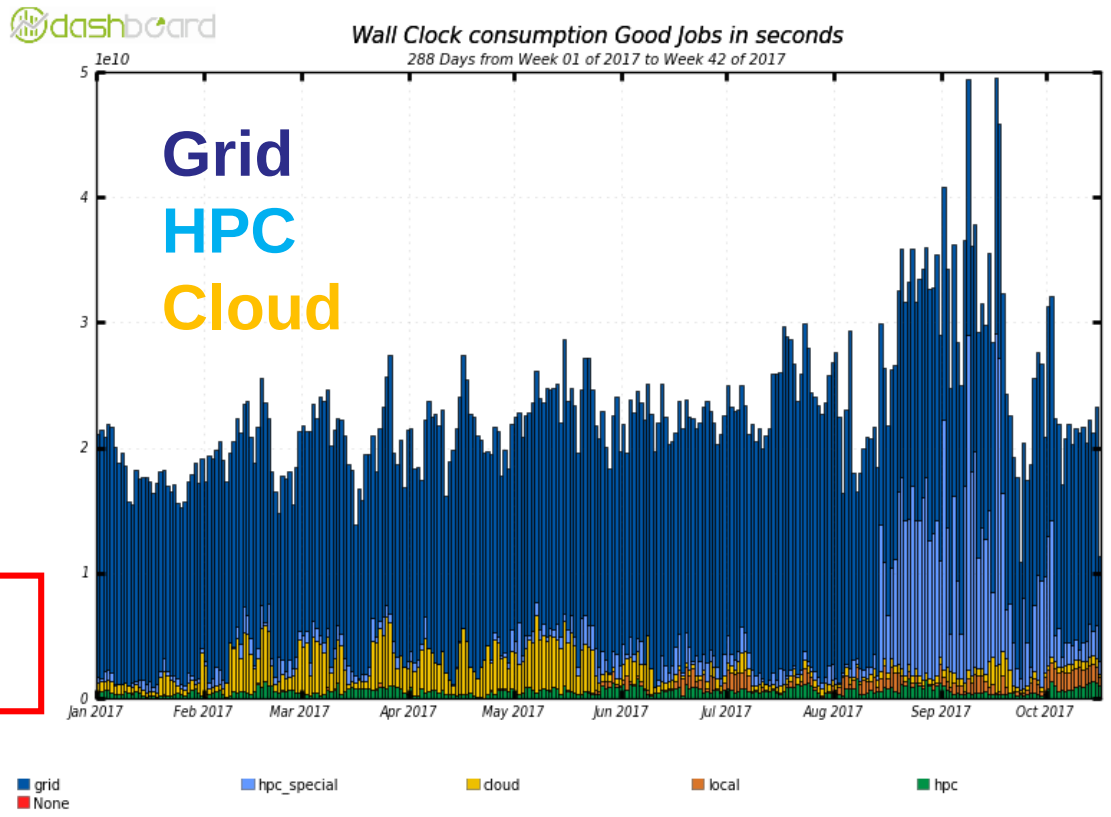

See Catherine, Eric, Wenjings's coming talks

Over 2017

- Grid 78%
- Cloud (academic, commercial, BOINC) 6.7%
- HPC_Special (NERSC_CORI, ORNL_Titan) 11%
- HPC (Local) 2.7%

Extra Resource (2)

- Cloud
  - Stable but no real increase
  - Cost gain wrt grid? Manpower?
- ATLAS@Home: Increasing
- HPC
  - Complex to set up
  - Test by CC@IDRIS. OK but 2.5k/10k slots max
- General
  - Harvester under development (common interface for ALL type of resource)
  - Event Service: Work at event level (simulation)

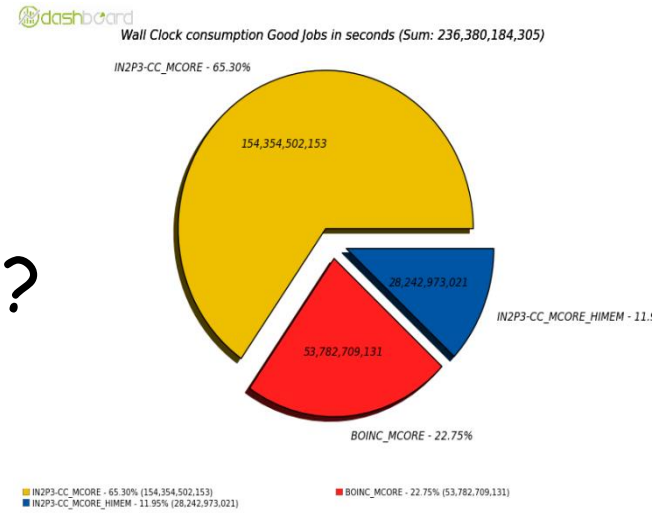

2017: WT BOINC (11k slots,22%) vs CCIN2P3

Data volume in 2017

1.2 EB!!

The Exabyte Area/Era
Dominated by Analysis

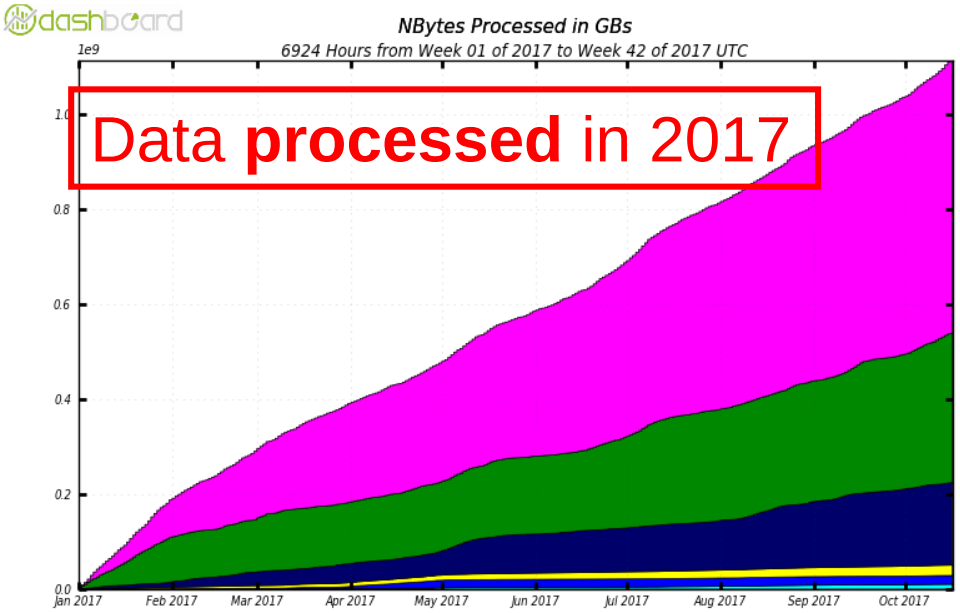

Data **processed** in 2017

70pB

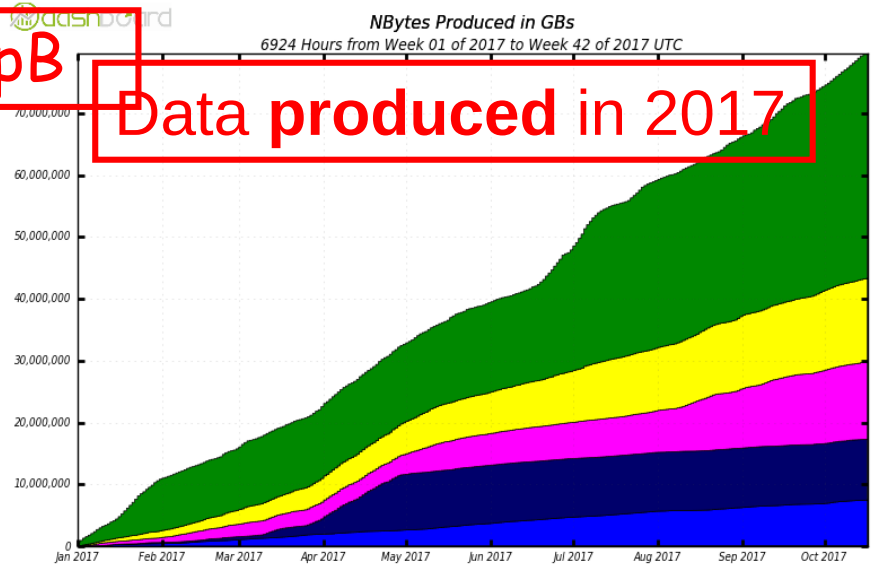

Data **produced** in 2017

942,141)
9,254)
))

Group Production (314,917,469)
Data Processing (18,547,935)
T0 Processing (0.00)
MC Reconstruction (173,950,675)
MC Simulation (10,493,575)

Total: 1,114,591,053 , Average Rate: 44.71 /s

Group Production (36,407,411)
MC Simulation (7,383,844)
Analysis (13,496,897)
Others (70,050)
MC Reconstruction (12,499,589)
unknown (0.00)
Data Processing (9,844,247)
T0 Processing (0.00)

Data Transfers (Rucio)

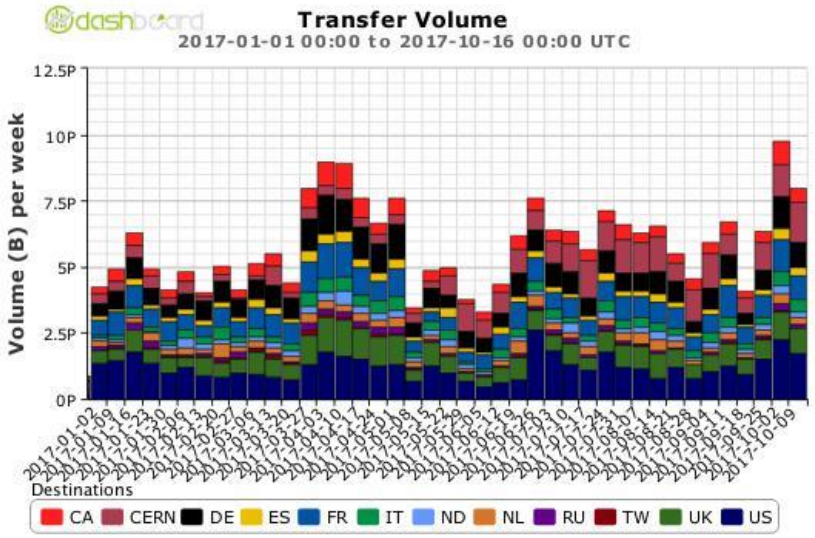

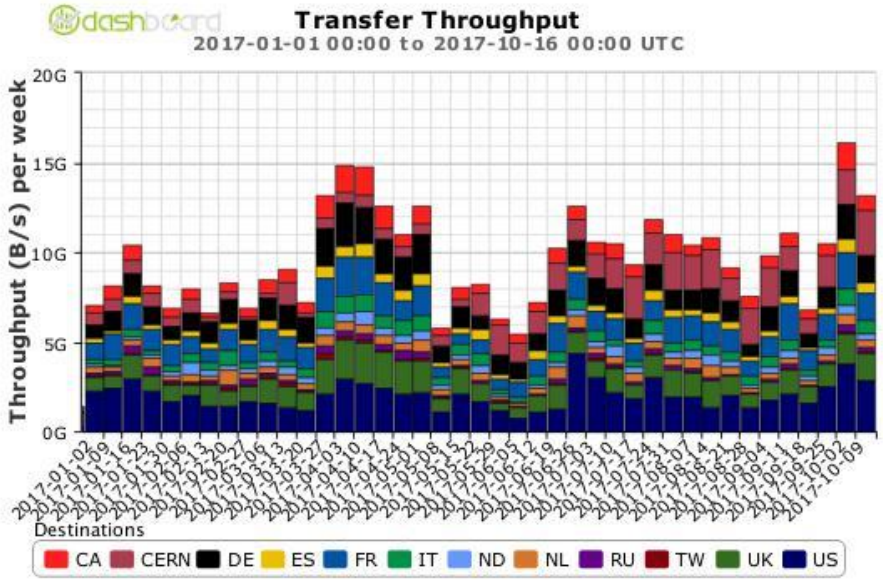

Per week:

- 7pB data transferred
- i.e. 15M files
- ~ 100Gb/s bandwidth
- 10pB deleted!

Automation helped:

- Data pre-placement ✓
- Data replication ✓
- Data rebalancing ✓

Finalizing Run-2

Feedback from Sites Jamboree (1)

- General https://indico.cern.ch/event/692124/
  - Some topics are ATLAS specific
  - Wider topics discussed at WLCG/HSF Workshop in Naples
- Few topics (my own prejudice)

- T2/T3 consolidation
- Event service for sites
- Protocols (SRM)
- CPU metrics
- CentOS7 migration
- Unified monitoring
- Containers/Singularity
- BOINC for sites
- Progress in monitoring
- Unified PQ
- WAN/LAN – Direct I/O
- Future: Data Lake/Ocean

Feedback from Sites Jamboree (2)

- T2/T3 consolidation
  - Policy endorsed by ICB
  - ATLAS recommends sites w/ small storage to focus in CPUs vs storage, ie become diskless
  - Limit
    - 0.5% ATLAS total storage, ie 400TB today, 600TB in 2020 (+15%/year)
  - In FR-cloud
    - Done in 2017 for RO-14 & RO-16
    - 2 sites below the limit Beijing (310TB) & RO-02 (345TB)
    - Start discussing?

Feedback from Sites Jamboree (3)

- Unified PQ
  - Single Panda queue for S/MCORE, LO & HIMEM
  - Includes Analysis, better handling of Score (eg Evgen)
  - OK for ARC (Cf. FZK), in progress for CREAM
- Event Service for sites
  - In progress and promising (opportunistic mode)
  - Some issues at large scale w/ brokering/priority
- Protocols
  - Goal: reach 90% sites w/ davs, root, (today 86%)
  - SRM still needed for tapes

Feedback from Sites Jamboree (4)

- WLAN vs LAN, Direct I/O vs CopytoScratch
  - Tests started (HC): seems OK for analysis, Prod?
  - ESS dev ongoing
- Towards HL-LHC
  - Real issue is storage.
    - Today shortage O(x10)
  - Progress on 'Data Lake' model
  - R&D with Google: 'Data Ocean'
    - Using Google cloud to store grid produced analysis outputs: 100% availability morning after big run
  - Core of discussion in Naples WLCG

I. Bird

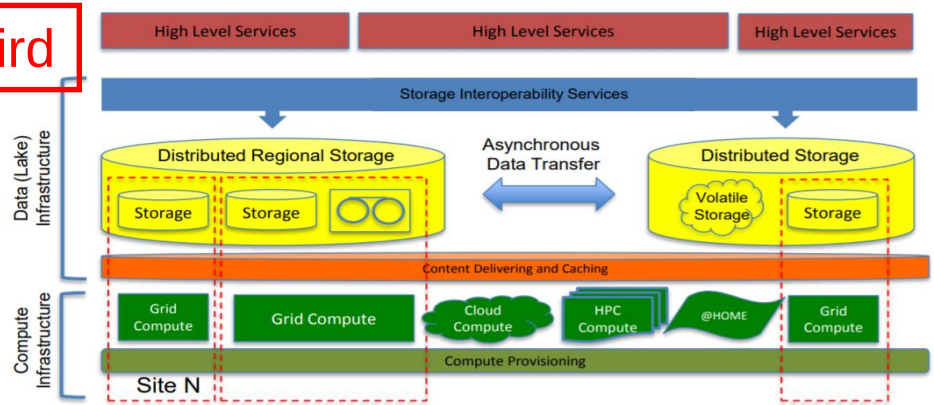

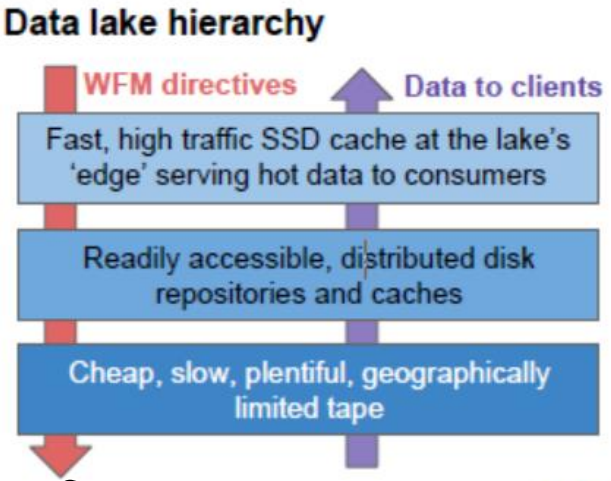

ICB feedback (1)

NB: My view as IN2P3 FA representative

- Hong-Kong T2 (HK-LCG2)
  - Got the green light to proceed as T2 (12/2017)!!
- Focus on Manpower
  - M&O effort on
    - ADC monitoring, Software performance (SPOT), Advanced workflow management (ES, ESS), Shifters
  - OTP
    - Class2: Missing DAST, ADCoS, CRC shifters
    - Class3: Missing 15% (Monitoring, Reco, Simu)
      - Too much Cloud squad coordination: 4.2FTEs
    - Class4: Reduction of needs in 2017

ICB feedback (2)

- Proposal for Manpower
  - FAs not yet providing CRC need to find someone
  - Cloud activities & uniformity
  - FAs should plan and invest in common projects
  - Filling ops oriented position (ATLAS funding)
- EuroHPC
  - New EU initiative
  - 2 phases: 2018-2020 (~200PF) & 2020-2024 (~1000PF) – 1.5B€ for Phase-1
  - WLCG not present yet!!
  - FA reps encouraged to discuss with EuroHPC reps how to proceed to bring some benefit to HEP

Additional Resource

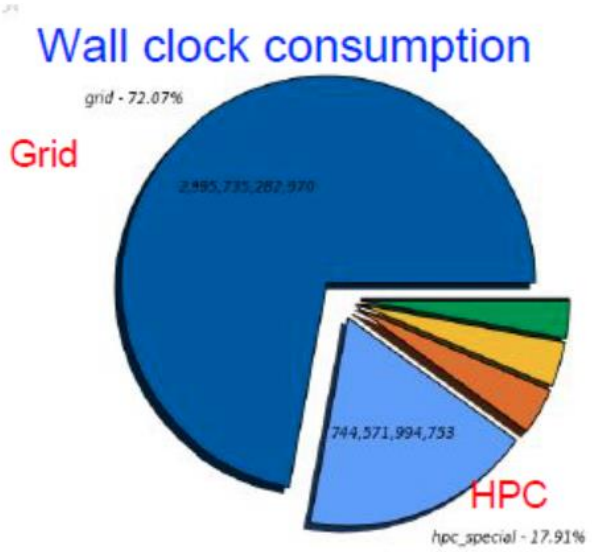

- ~20% from HPC
- Titan reached 96% efficiency with ATLAS backfill
- 17% of simulated events delivered by Event Service

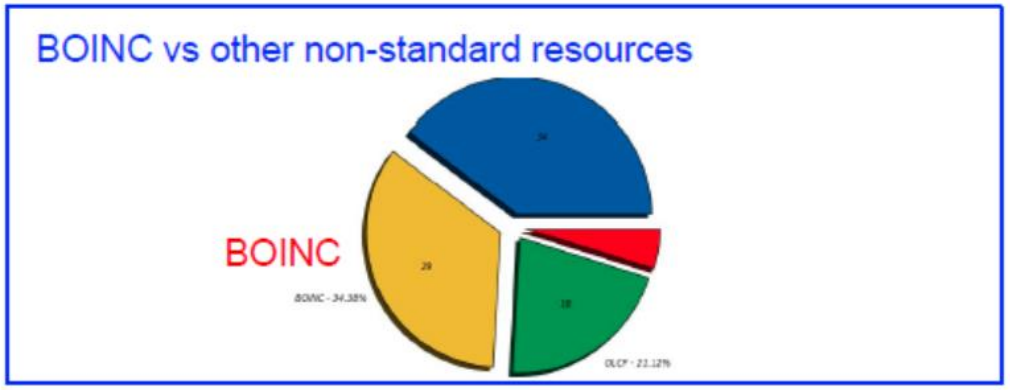

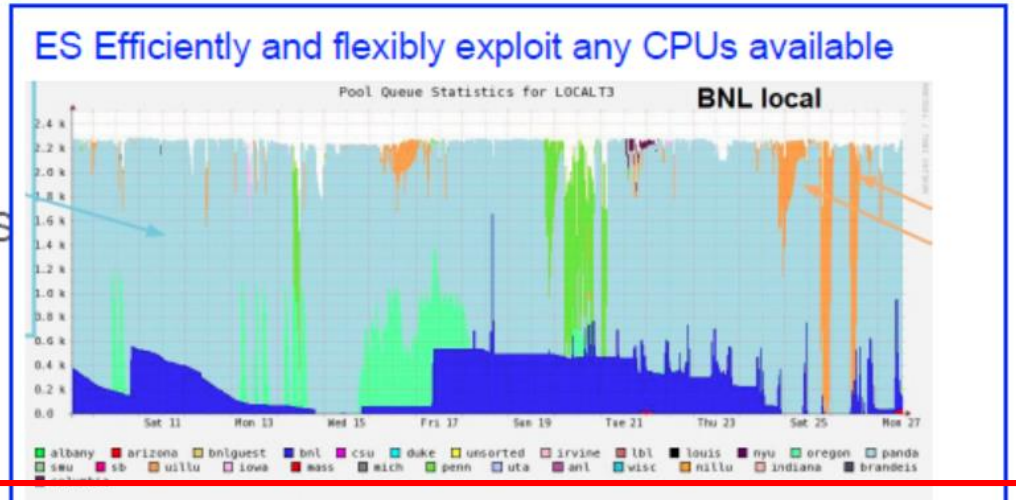

Event Service, Event Streaming Service, Harvester

- What ES is to computing, ESS is to input data transfer
  - Instead of transferring whole input file, transmit only the range of evts needed

US-HPC (in March)

Running slots (Simulation only)

Mostly NERSC (Berkeley)

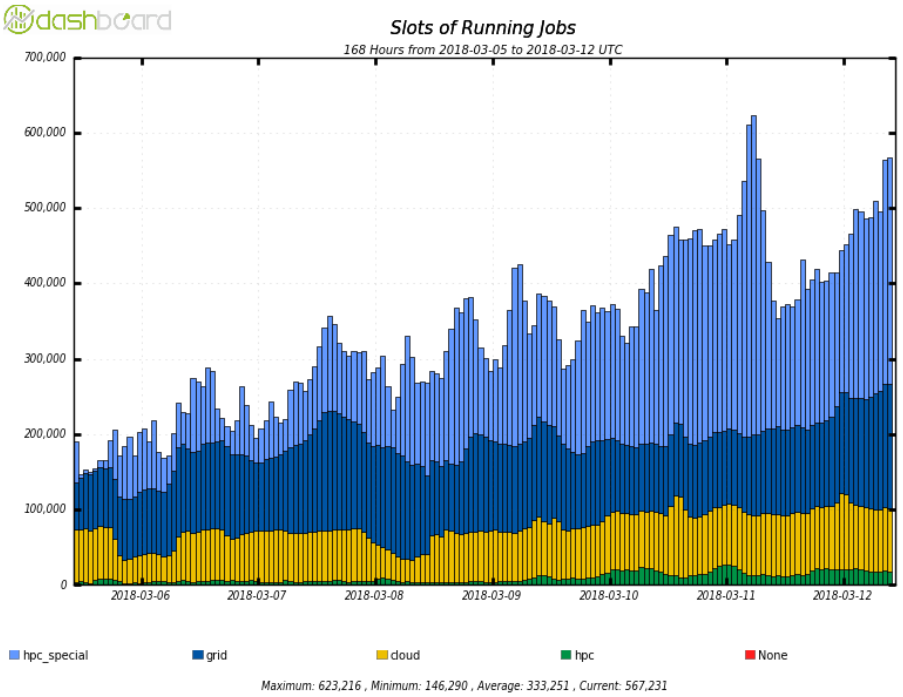

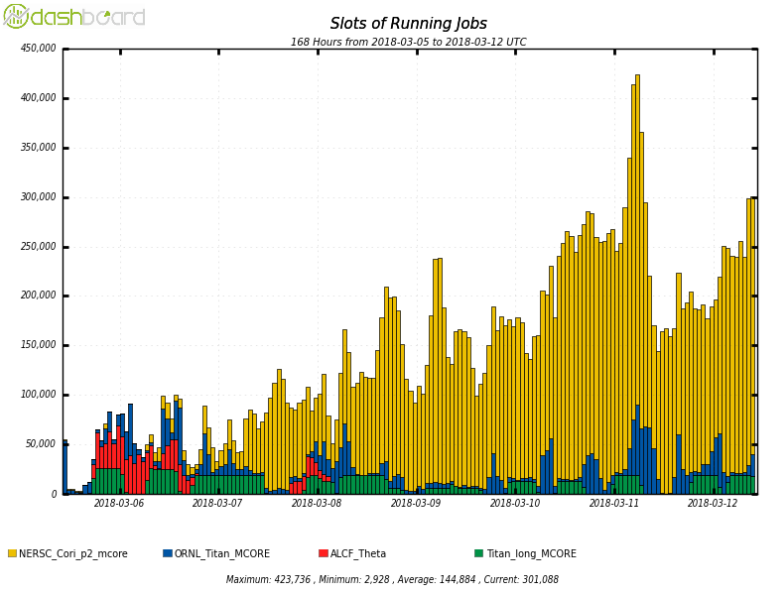

Huge Contribution last week from US-HPC

- Up to 400K slots
- **NB: HS06 (HPC) ~ 0.1 * HS06 (Grid)**

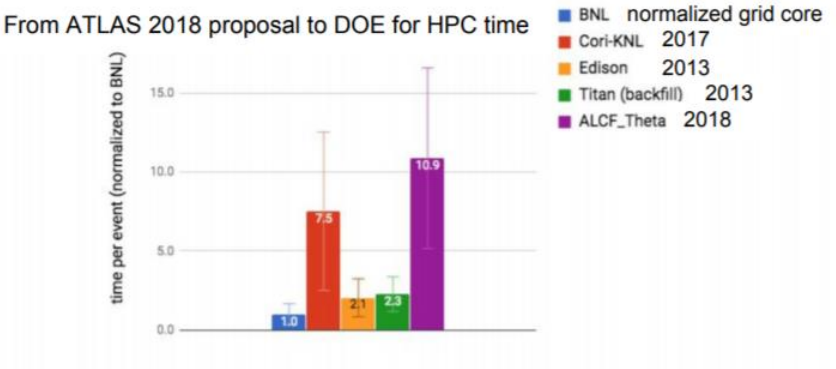

Lines of effort (1)

- From Physics Wkshop with 120fb^{-1} 12/2017: Limited Monte Carlo statistics!
  - Leverage additional resources (HPC, Boinc, …)
  - Improve software and efficiency (SPOT)
    - SPOT: Software performance optimisation team
    - Technical optimisation could already provide 20% CPU gain CPU time (eg Compression)
  - Run less full-simulation (and more fast sim)

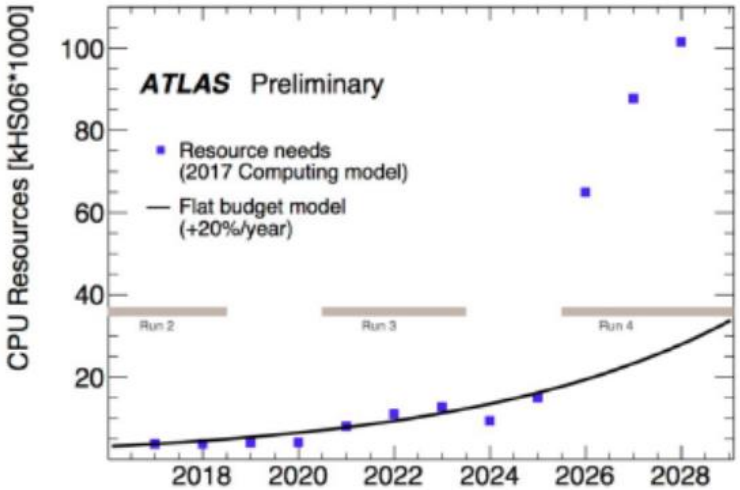

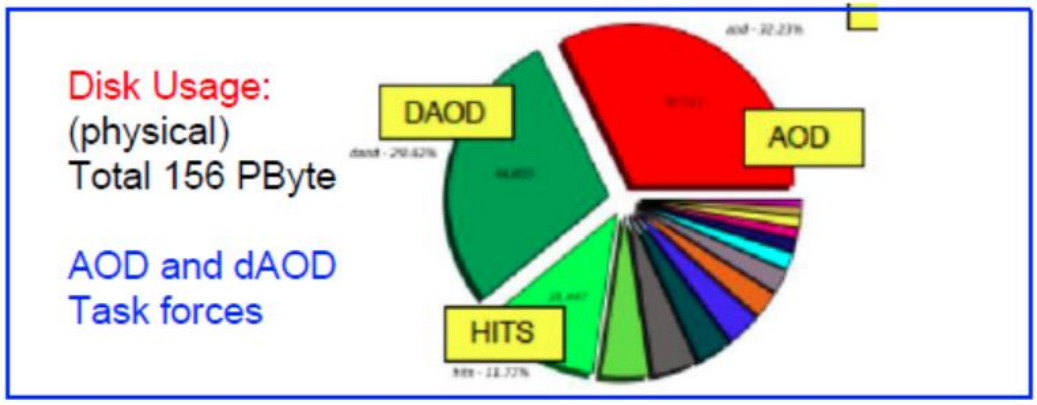

Lines of effort (2)

- T1s continue to exercise and improve perf. of DAOD production from tape inputs
- Promote support for software development
  - Supporting software activity where the effort is crucial and currently insufficient
- Support the roll-out of containers
  - Powerful technology for improving uniformity, ease of ops & security across resources. Singularity deployment model agreed by WLCG and ATLAS. Proactive deployment on grid sites is proceeding
- T3s policy: Will to limit 'bad' T3s / small sites

Preparing Run-3

Pledges 2019 (prel.)

- C-RSG April: No changes to 2019 requests Oct->now
- Preliminary questions from C-RSG
  - HI run plans, coping w/ +20% lumi, mcore eff., tape usage,T0 spillover impact to grid, tape based workflows (prestaging 'data carousel')
- Preliminary pledges (from Rebus)

	2018 (k)	2019 (k)	Delta (%)
T0 CPU	411	411	0
T0 Disk	26	27	4
T0 Tape	94	105	12
T1 CPU	949	1057	11
T1 Disk	72	88	22
T1 Tape	195	221	13
T2 CPU	1160	1292	11
T2 Disk	88	108	23

- Modest increase
- Emphasis on Disk @ T1&T2
- Inside flat budget (except disk)
- Not final/approved yet!!

Progress on simulation

- Dominates CPU needs and adds systematic to some physics studies (eg VHbb)
- FastSimulation
  - Validated in Sep'2017
  - Uses FastCaloSim
  - Test of FastChain
- Pile-up treatment
  - Standard vs Overlay

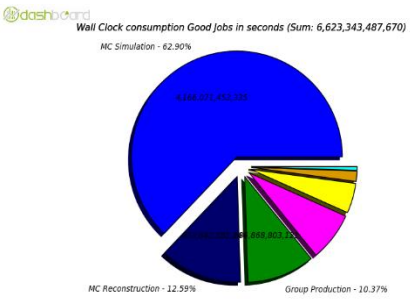

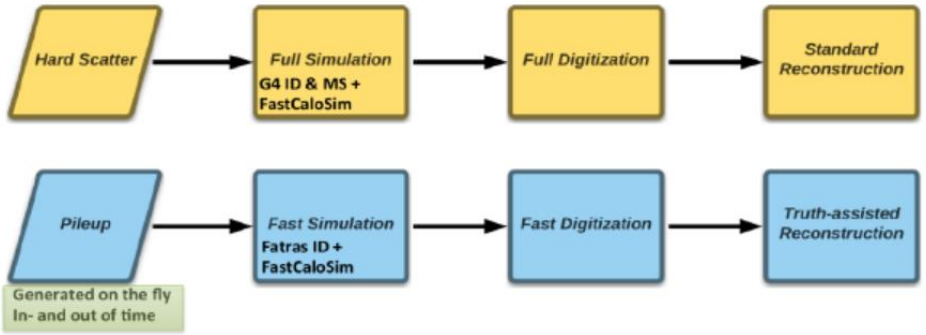

- Combining multiple HITS level events during pile-up digitization is not efficient in terms of CPU or I/O.
  - Event Digitization time w/pile-up ~ Event Simulation time w/ATLFASTII
  - Current approach relies on high I/O to avoid high memory usage.
  - The situation will get worse as <mu> increases.
- Data Overlay (combining a digitized hard-scatter event with a Zero-bias data event) has been used successfully for heavy ion campaigns.
  - Lower CPU, memory and I/O requirements.
  - Places a huge load on conditions database infrastructure, due to jobs requesting different conditions IoVs.

Analysis: Issues & Progress

- Run2 Model (xAOD, Train & Derivations) OK
- Assumed
  - Derivations run fast enough -> able to run new full derivations every few weeks if necessary
  - Derivations output small -> grid analysis jobs able to process all data/MC in O(1 day)
- In real
  - ~6 weeks for repro all data into O(80) DAOD
    - -> Run only major productions
  - Analysis jobs tails >>1 day
    - Write larger outputs to reduce #grid iterations
- Solutions
  - DAOD production: More efficient merging, reduce size
  - DAOD processing: Improvement in Distributed Analysis

Towards Run-3

- Going slowly away from hierarchical T0/T1/T2 mode: Nucleus/satellites model
  - Storage, transfers & resource optimized
    - eg Now T2s disk are full
  - I/O further optimized
- New Database scheme
  - To replace COOL & handle Overlay (IRFU, LAL)
- Key words
  - Harvester, Event Service, Overlay
  - Containers: Virtualization for batch execution (Singularity)

Reconstruction for Run-3

- Going from multi-event parallelism (MP) to inter-event parallelism @algo level (MT)
- Compulsory for high-pileup tracking (Run-3)
  - More memory needed for Reco (pile-up)
  - Memory/core available not increasing
- Ready for Run-3 start
  - Algorithmic migration will take some time

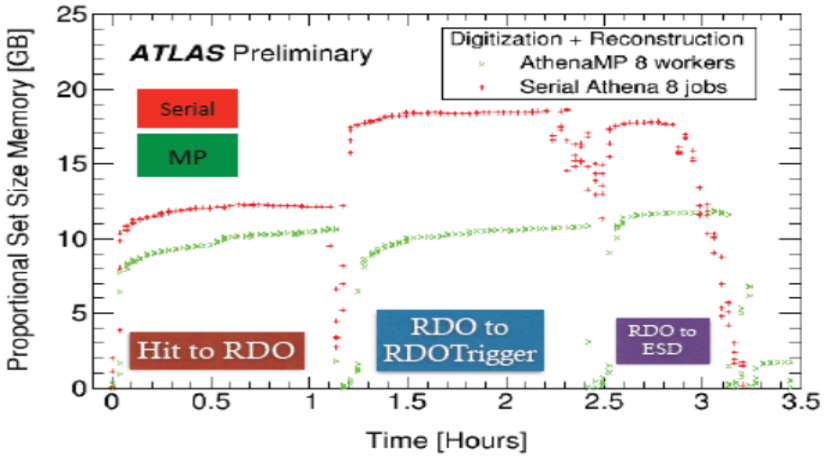

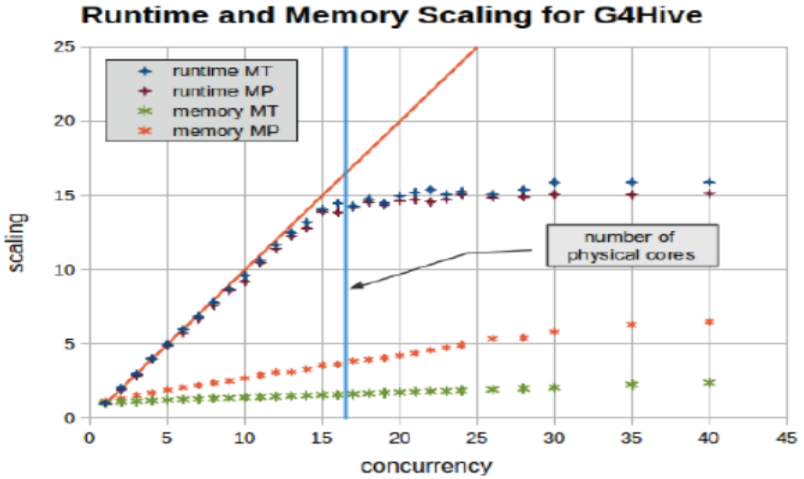

Towards HL-LHC

Towards HL-LHC: Inputs

T. Wenaus

Input Parameters at HL-LHC,
updated after the conclusion of the Layout Task Force

Output HLT rate: 10kHz
Reco time: 130s/event at mu=200, Simul Time: 454 s/event
Nr Events MC / Nr Events Data = 1.5
N events with Fast Simulation: 50% of Full Simulation
LHC live seconds /year: 7.3 M

Flat budget
20% more CPU/year, 15% more storage/year
Evolution from a 2017 baseline
Data from previous years taken into account
Tier-0 contribution added to the total

HEP Software Foundation (HSF)

- Goal: facilitates coordination & common efforts in HEP software and computing internationally
- ATLAS part of it & benefit from HSF
  - In WFM and DDM we have strong software we should advance as possible community standards
    - Rucio, PanDA, Harvester, the package of all of them
- Output: CWP (Community White Paper)
  - Scope: HL-LHC
  - Items addressed: Flexible management of facilities, use of heterogeneous resource, Computing Models, Facilities, Distributed Computing
  - Made public 12/2017
  - http://hepsoftwarefoundation.org/index.html

CWP, R&D, Common solutions

- HSF CWP process completing -> focus to identified R&D topics and common solutions
  - http://hepsoftwarefoundation.org/index.html
  - Well positioned with hardened, scalable systems
    - First Rucio Community Workshop coming in March
    - PanDA adoption beyond ATLAS ongoing, e.g. Compass
  - Good start on R&D efforts towards Run 3 & 4
    - Event Service produced 20% of all ATLAS MC last week
    - Event streaming service R&D underway, planning data lake R&D
  - Discussions of R&D and common solutions at WLCG 2018 & CHEP 2018

Towards HL-LHC: CPU

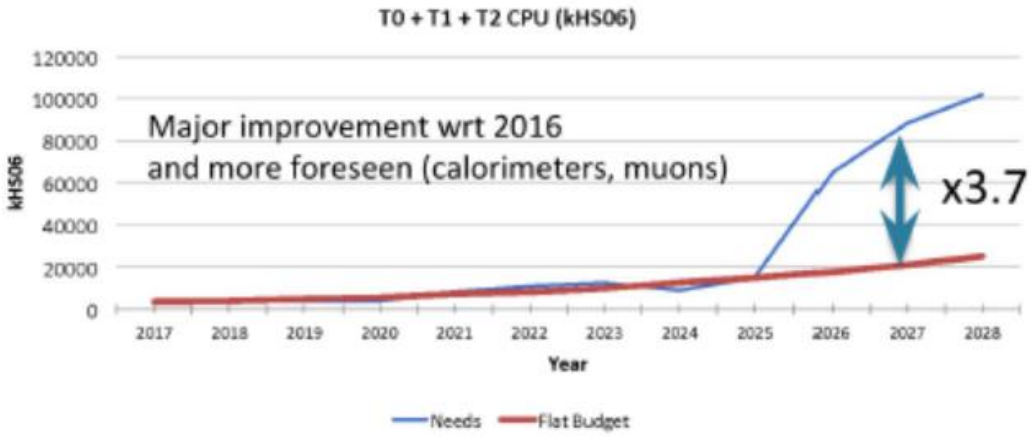

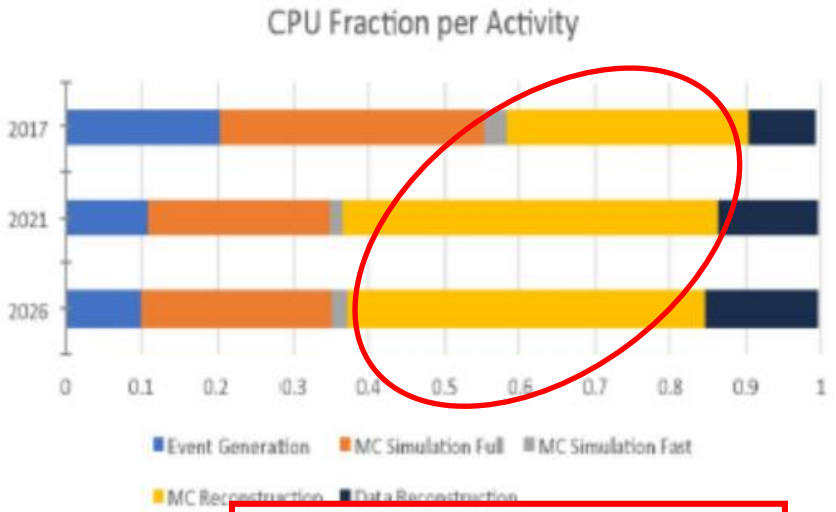

Dominated by MC Reco

CPU: the gap has shrunk significantly with e.g. substantial improvements in reco time and improvements of the model

As other subsystems follow ITk in making reco improvements we can expect the gap to shrink further

Also relies on a much larger role for fast simulation

Towards HL-LHC: Disk

T. Wenaus

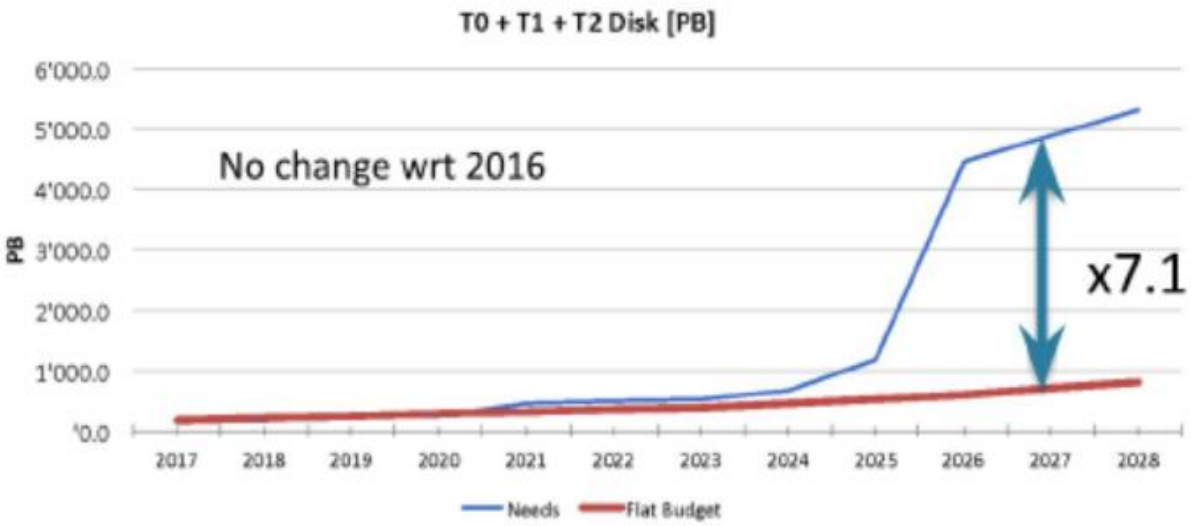

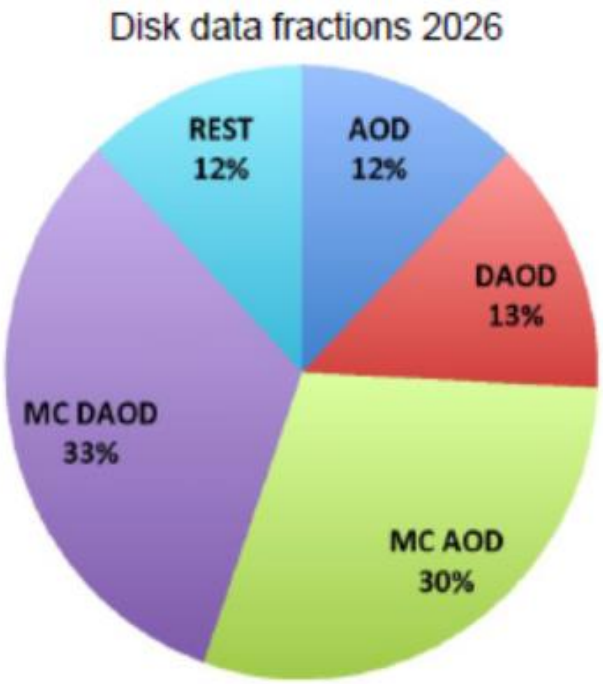

Storage: the gap won't shrink until we develop and quantify the strategies to bring it down

This we are starting to do, e.g. with a program of testing and improving train workflows using tape as input

Possible gains

- Improvement on CPU
  - Detector layout (TDR)
  - Machine learning technique
  - Fast simulation/Fast chain
- Improvement in Storage
  - No AOD on disk (Run Train analysis from AOD on tape
- Not enough. Gain to come
  - From re-thinking of distributed storage an& data access
  - A network driven data model allows to reduce the amount of storage, eg disk

Network driven 'data lake'

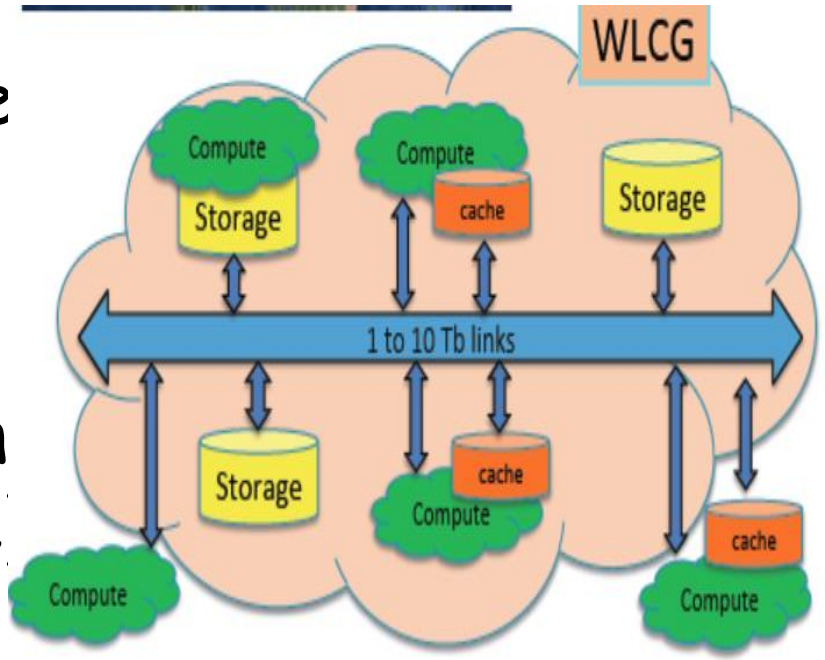

Summary

- ATLAS computing in very good shape!
- Now able to focus on refinements, performance, and look to future with R&D
- ATLAS should be front and center in common R&D (inside HSF community)
- Run-3 a priori OK within flat budget. Key issue is software: AthenaMT
- HL-LHC
  - Trend lines are good in CPU (constant progress)
  - Plans in storage to be quantified (today critical)
  - R&D, 'Data lake' model, ~~non-flat budget?~~