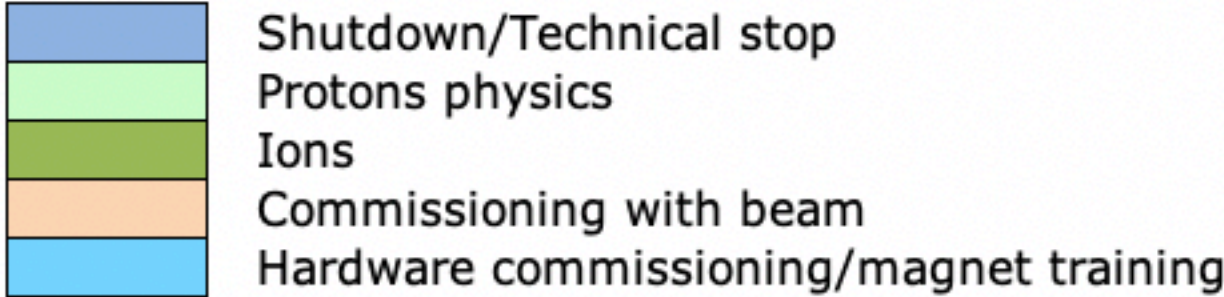



# ATLAS: Optimising ATLAS throughput on GridPP

J. Walder

GridPP46, Ambleside - September 2021

# Timeline



Last updated: June 2021

Mode	GPDs	LHCb	ALICE
p-p	160/fb	25-30/fb (~50/fb by LS4)	200/pb

<https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm>



# Computing throughput on GridPP

- UK average of 483k HS06 over last year.

- 430k (for first 3 months),

- 520k (last 3 months)



- Peaks nowrunning above 600k HS06

- Major activities:

- MC reco.
- MC evgen.
- Simulation
- Group prod.
- User analysis


Federation	Avg. HS06	2020 pledge	2021 pledge
UK-T1-RAL	~156 K	156k	173k
UK-NorthGrid	136 K	51k	50k
UK-ScotGrid	91.9 K	33k	45k
UK-London-Tier2	79.4 K	40k	43k
UK-SouthGrid	37.9 K	21k	22k

Activity	Avg. HS06
MC Reconstruction	119 K
MC Event Generation	97.7 K
MC Simulation Full	89.0 K
Group Production	84.6 K
User Analysis	45.9 K
MC Simulation Fast	12.2 K
Data Processing	11.6 K
Group Analysis	8.31 K
MC Resimulation	6.77 K
MC Merge	1.24 K



# UK Contributions


- Relative to T1/2 at other clouds (excl. Cern) in last 12 months:
  - ~13% of completed jobs and ~10% of running slots.




- Very similar job success rates cf whole T1/T2 sites
- Main failure modes relate to stage-in / out failures or failures during direct-IO reading

## Job completion status

### ALL T1/T2

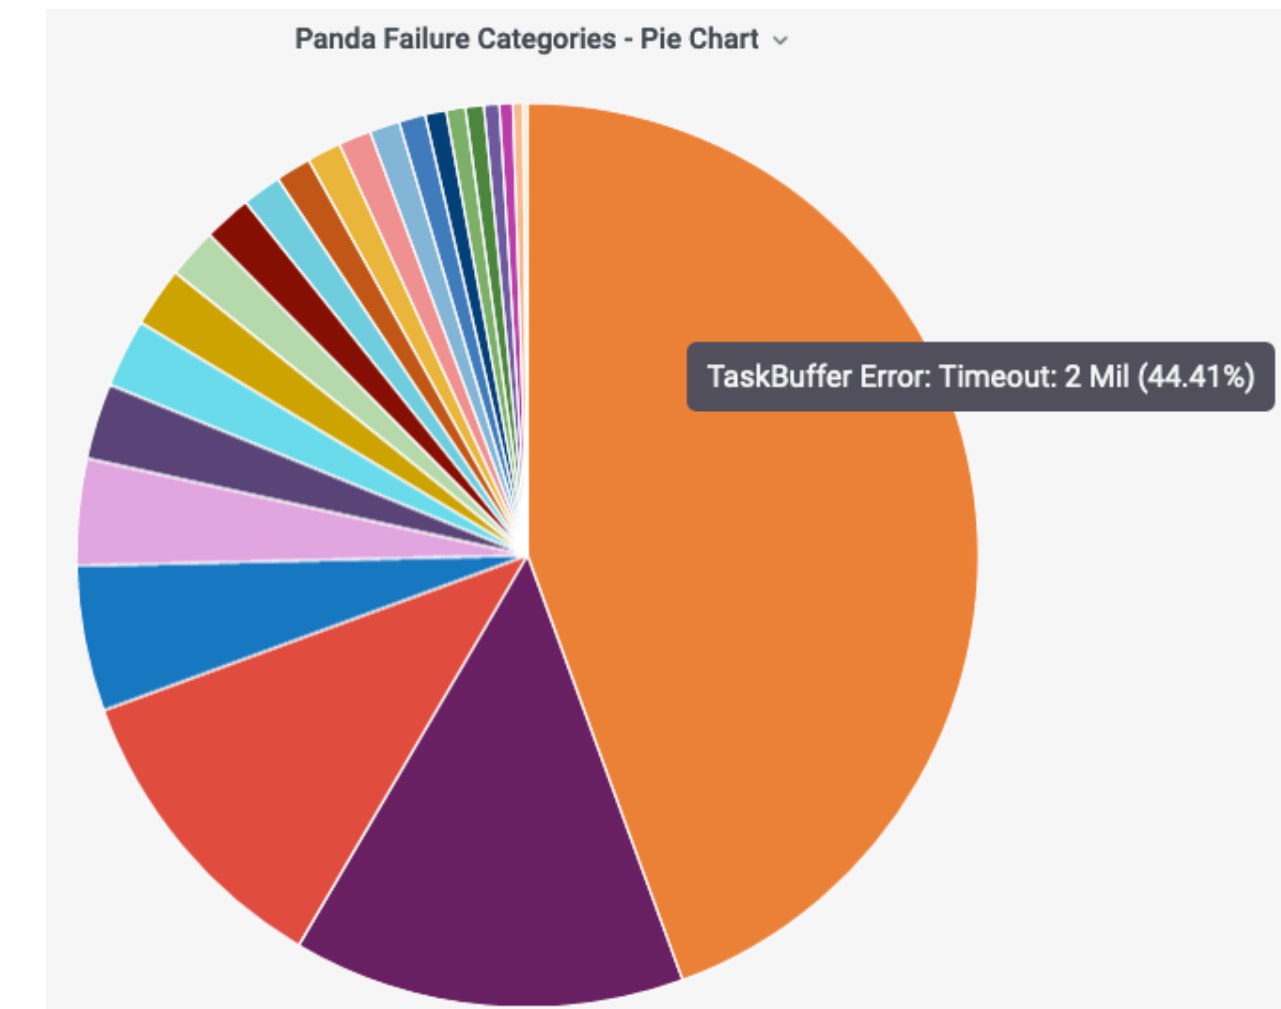


### UK T1/T2



## Job failure modes

### UK T1/T2





	total	percentage
TaskBuffer Error: Timeout	2 Mil	44%
Pilot Error	589 K	14%
TaskBuffer Error 300	460 K	11%
TaskBuffer Error	218 K	5%
Athena/Exec Error: Athena	160 K	4%
Job Dispatcher Error: lost heartbeat	114 K	3%
Execution Error 65	102 K	2%
Pilot/PanDA Error: Get error	90 K	2%
Job Dispatcher Error	75 K	2%
Pilot/DDM Error: get error	72 K	2%
Transformation Error	59 K	1%
Pilot/DDM Error: put error	54 K	1%
Transformation Error: not installed in CE	51 K	1%
DDM Error	49 K	1%
Athena/Exec Error: Proot	45 K	1%
Pilot/PanDA Error	40 K	1%
Pilot/DDM Error	32 K	1%
UNKNOWN	29 K	1%
Pilot/PanDA Error: Put error	28 K	1%




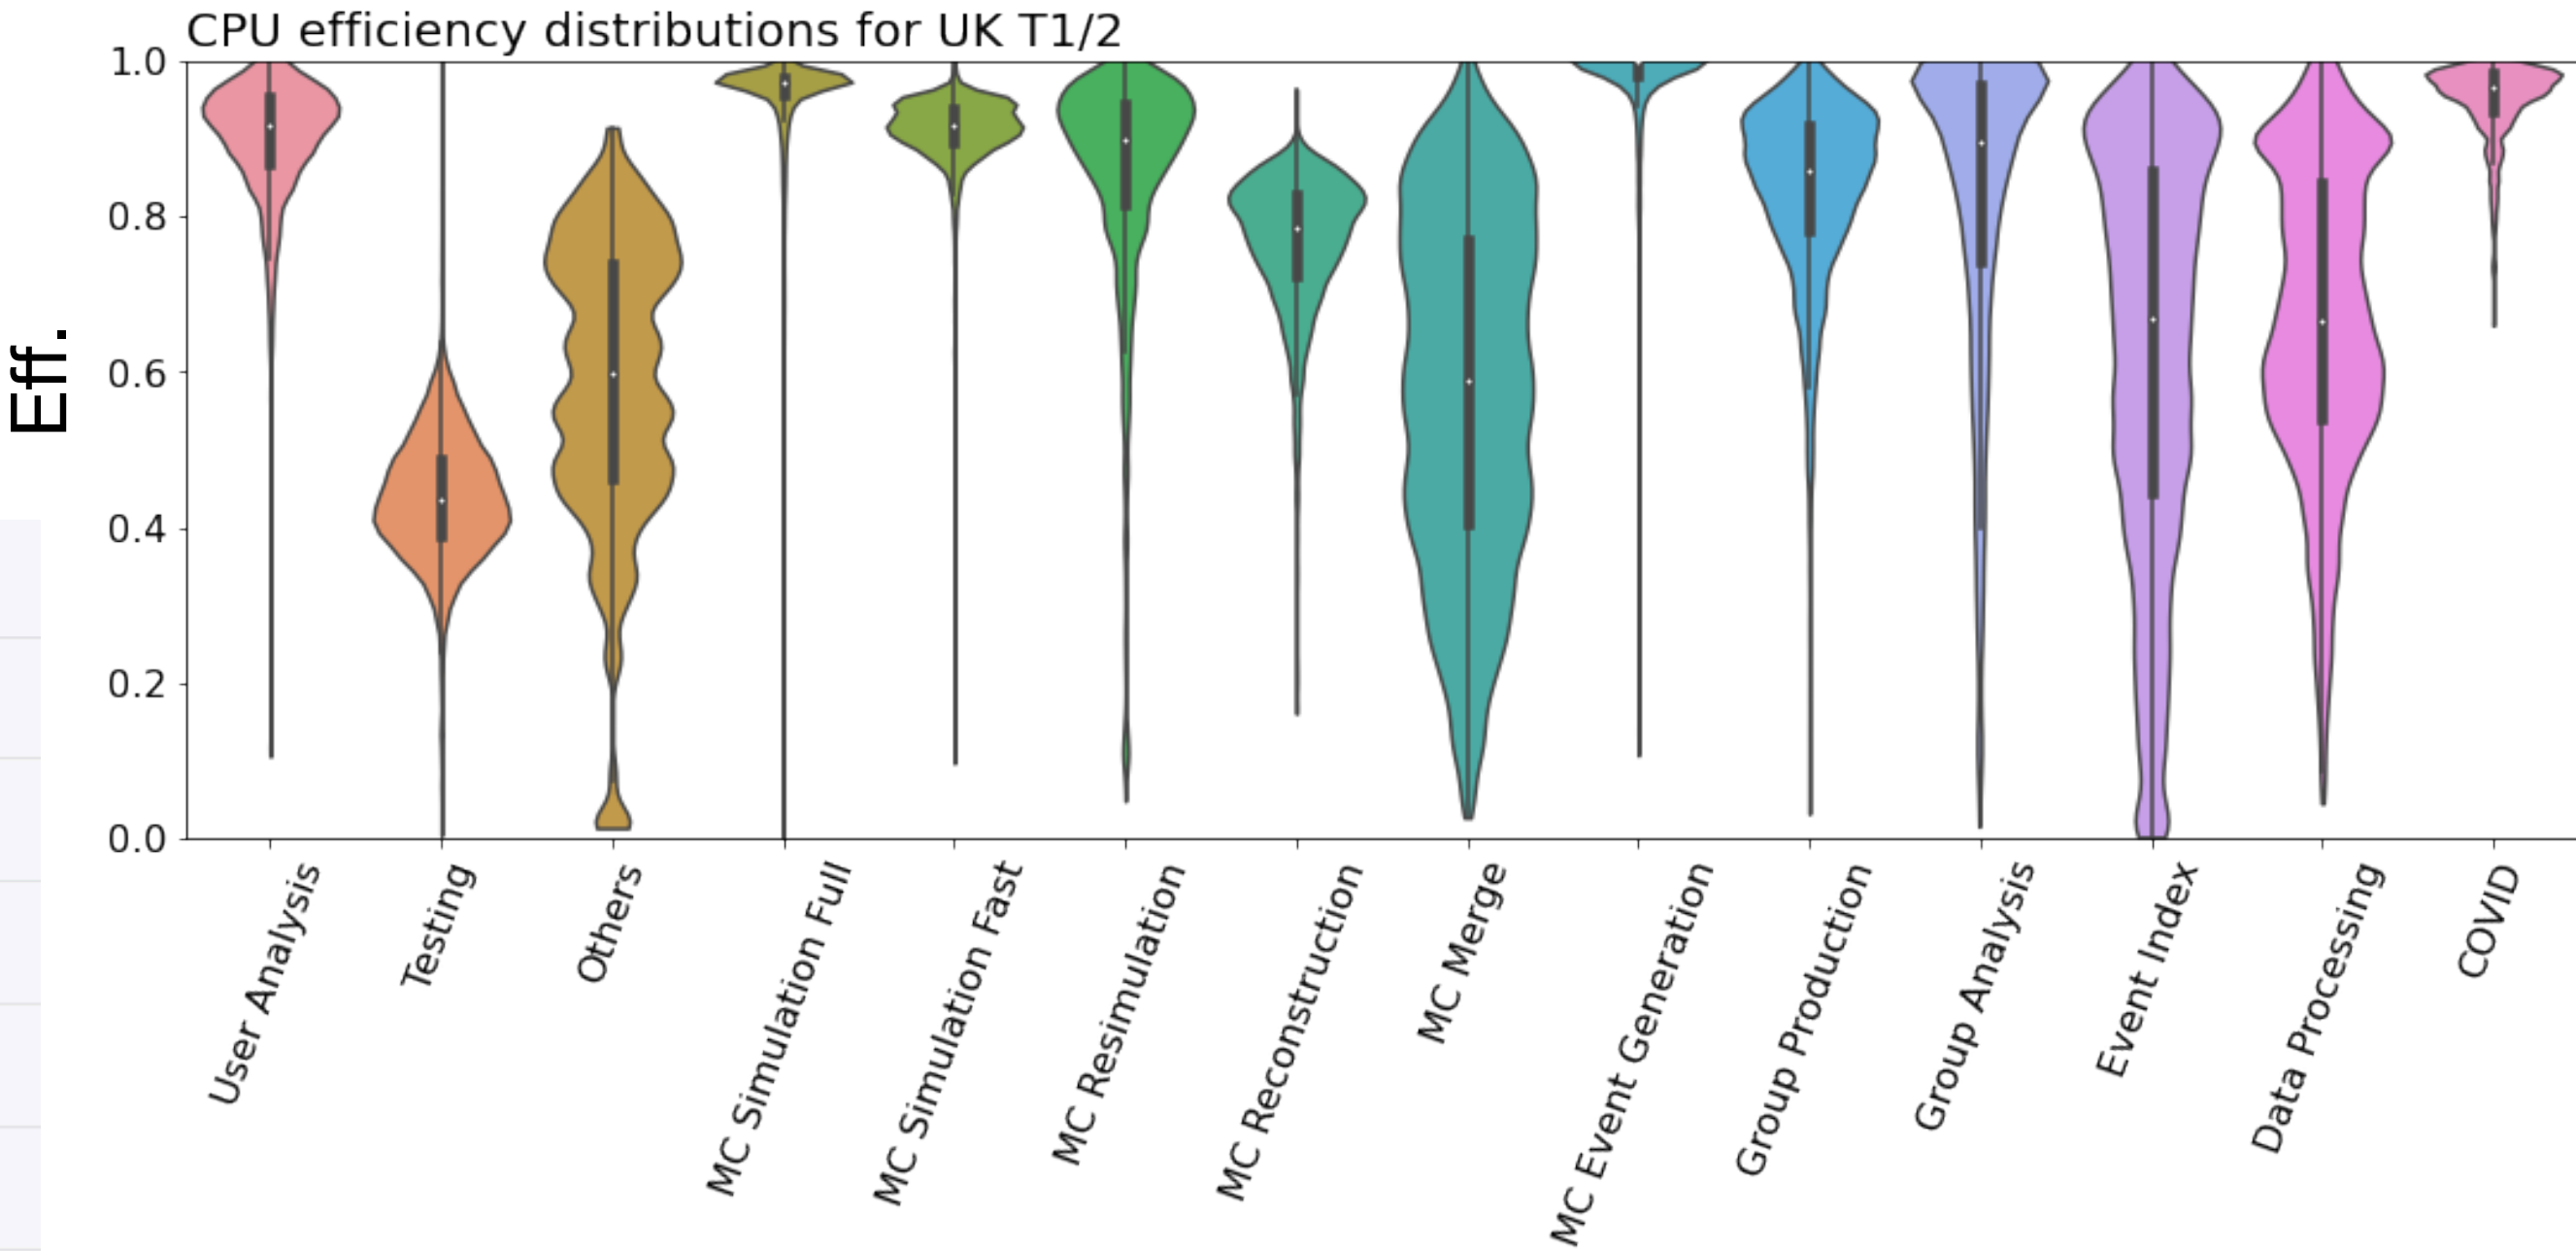
# CPU Efficiencies

- Broad range of CPU efficiencies for various job types:
- ~ 60% average CPU efficiency for IO intensive (e.g. MC merge) computations for UK sites
  - Usual average metrics hide broad distributions in many job types

CPU efficiencies at UK (T1/2) sites for given job types in last year




Distributions (hourly averaged) of underlying data




# GridPP Storage

- Site -> site transfers
  - ALL (excl. UK) -> UK
    - 63 PB transfers
  - UK->ALL (excl. UK)
    - 58 PB transfers
  - UK->UK
    - 12 PB transfers

Production Input	22.8 PB
Analysis Input	15.9 PB
Data Consolidation	14.6 PB
Production Output	8.76 PB
Staging	7.27 PB
User Subscriptions	3.74 PB



- Intra-site transfers (or from remote SE for storageless sites) for inputs for batch Jobs:
  - 270 PB data transferred in.






- Regular deletion campaigns performed
  - Obsolete datasets (e.g superseded by more recent processings)
    - Analysers request exceptions for needed datasets.
  - Lifetime-based removal of untouched datasets
- 90PB of Deletions (over 150M objects) during this period.

Federation	Disk	Pledge
UK-London-Tier2	4635 TB	4046 TB
UK-NorthGrid	7797 TB	6730 TB
UK-ScotGrid	3946 TB	3094 TB
UK-SouthGrid	691 TB	690 TB
UK-T1-RAL	15540 TB	15540 TB

# Other milestones

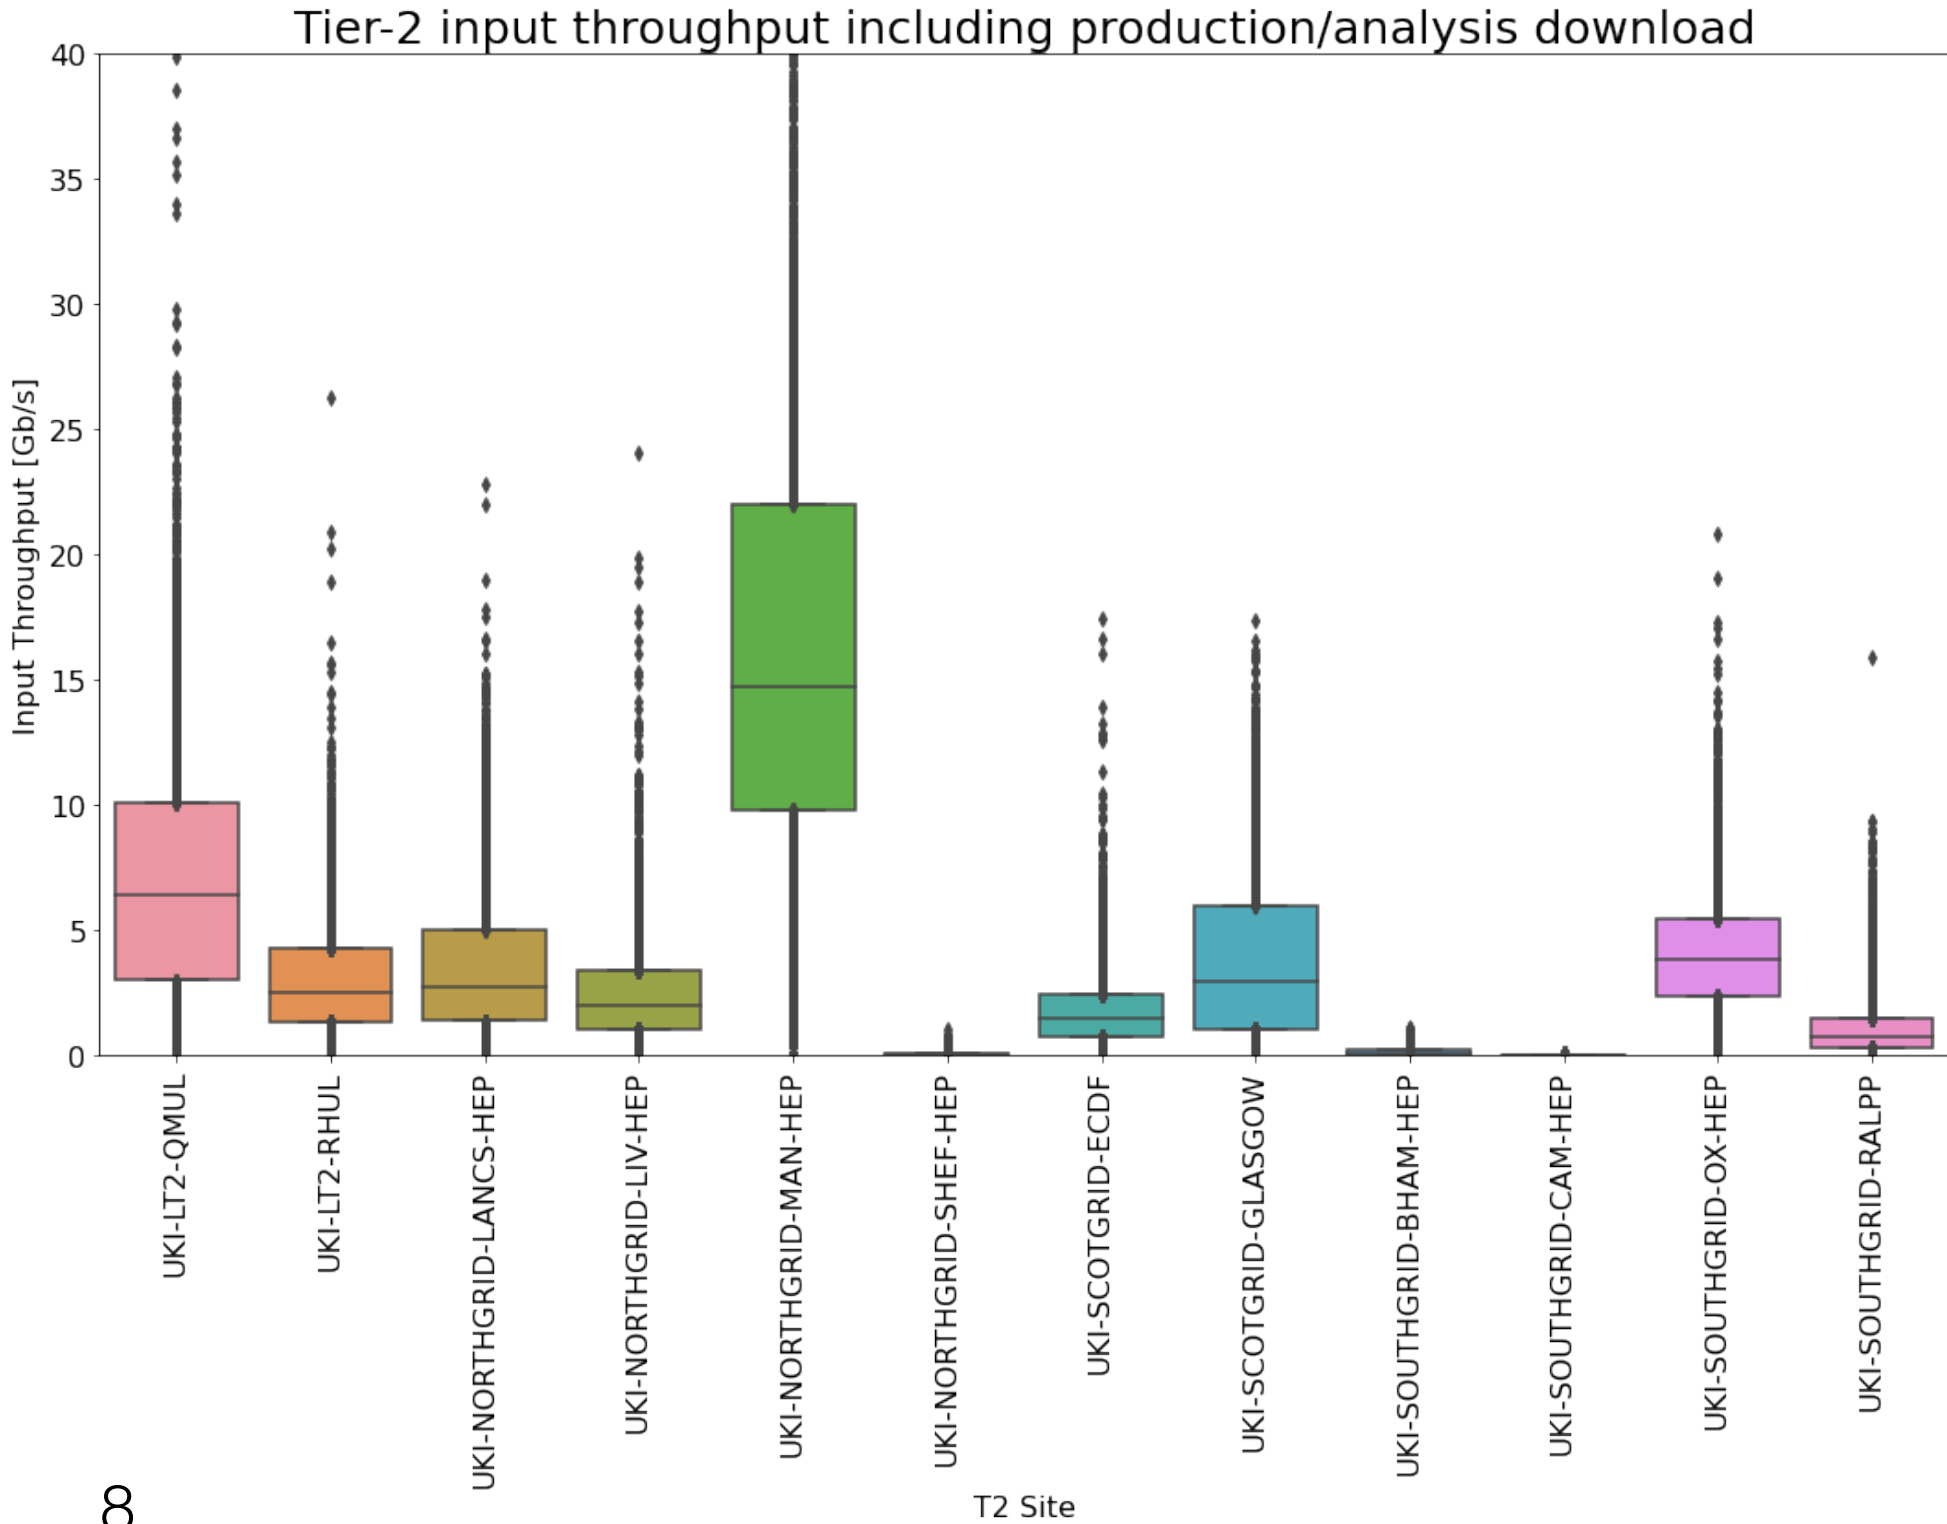
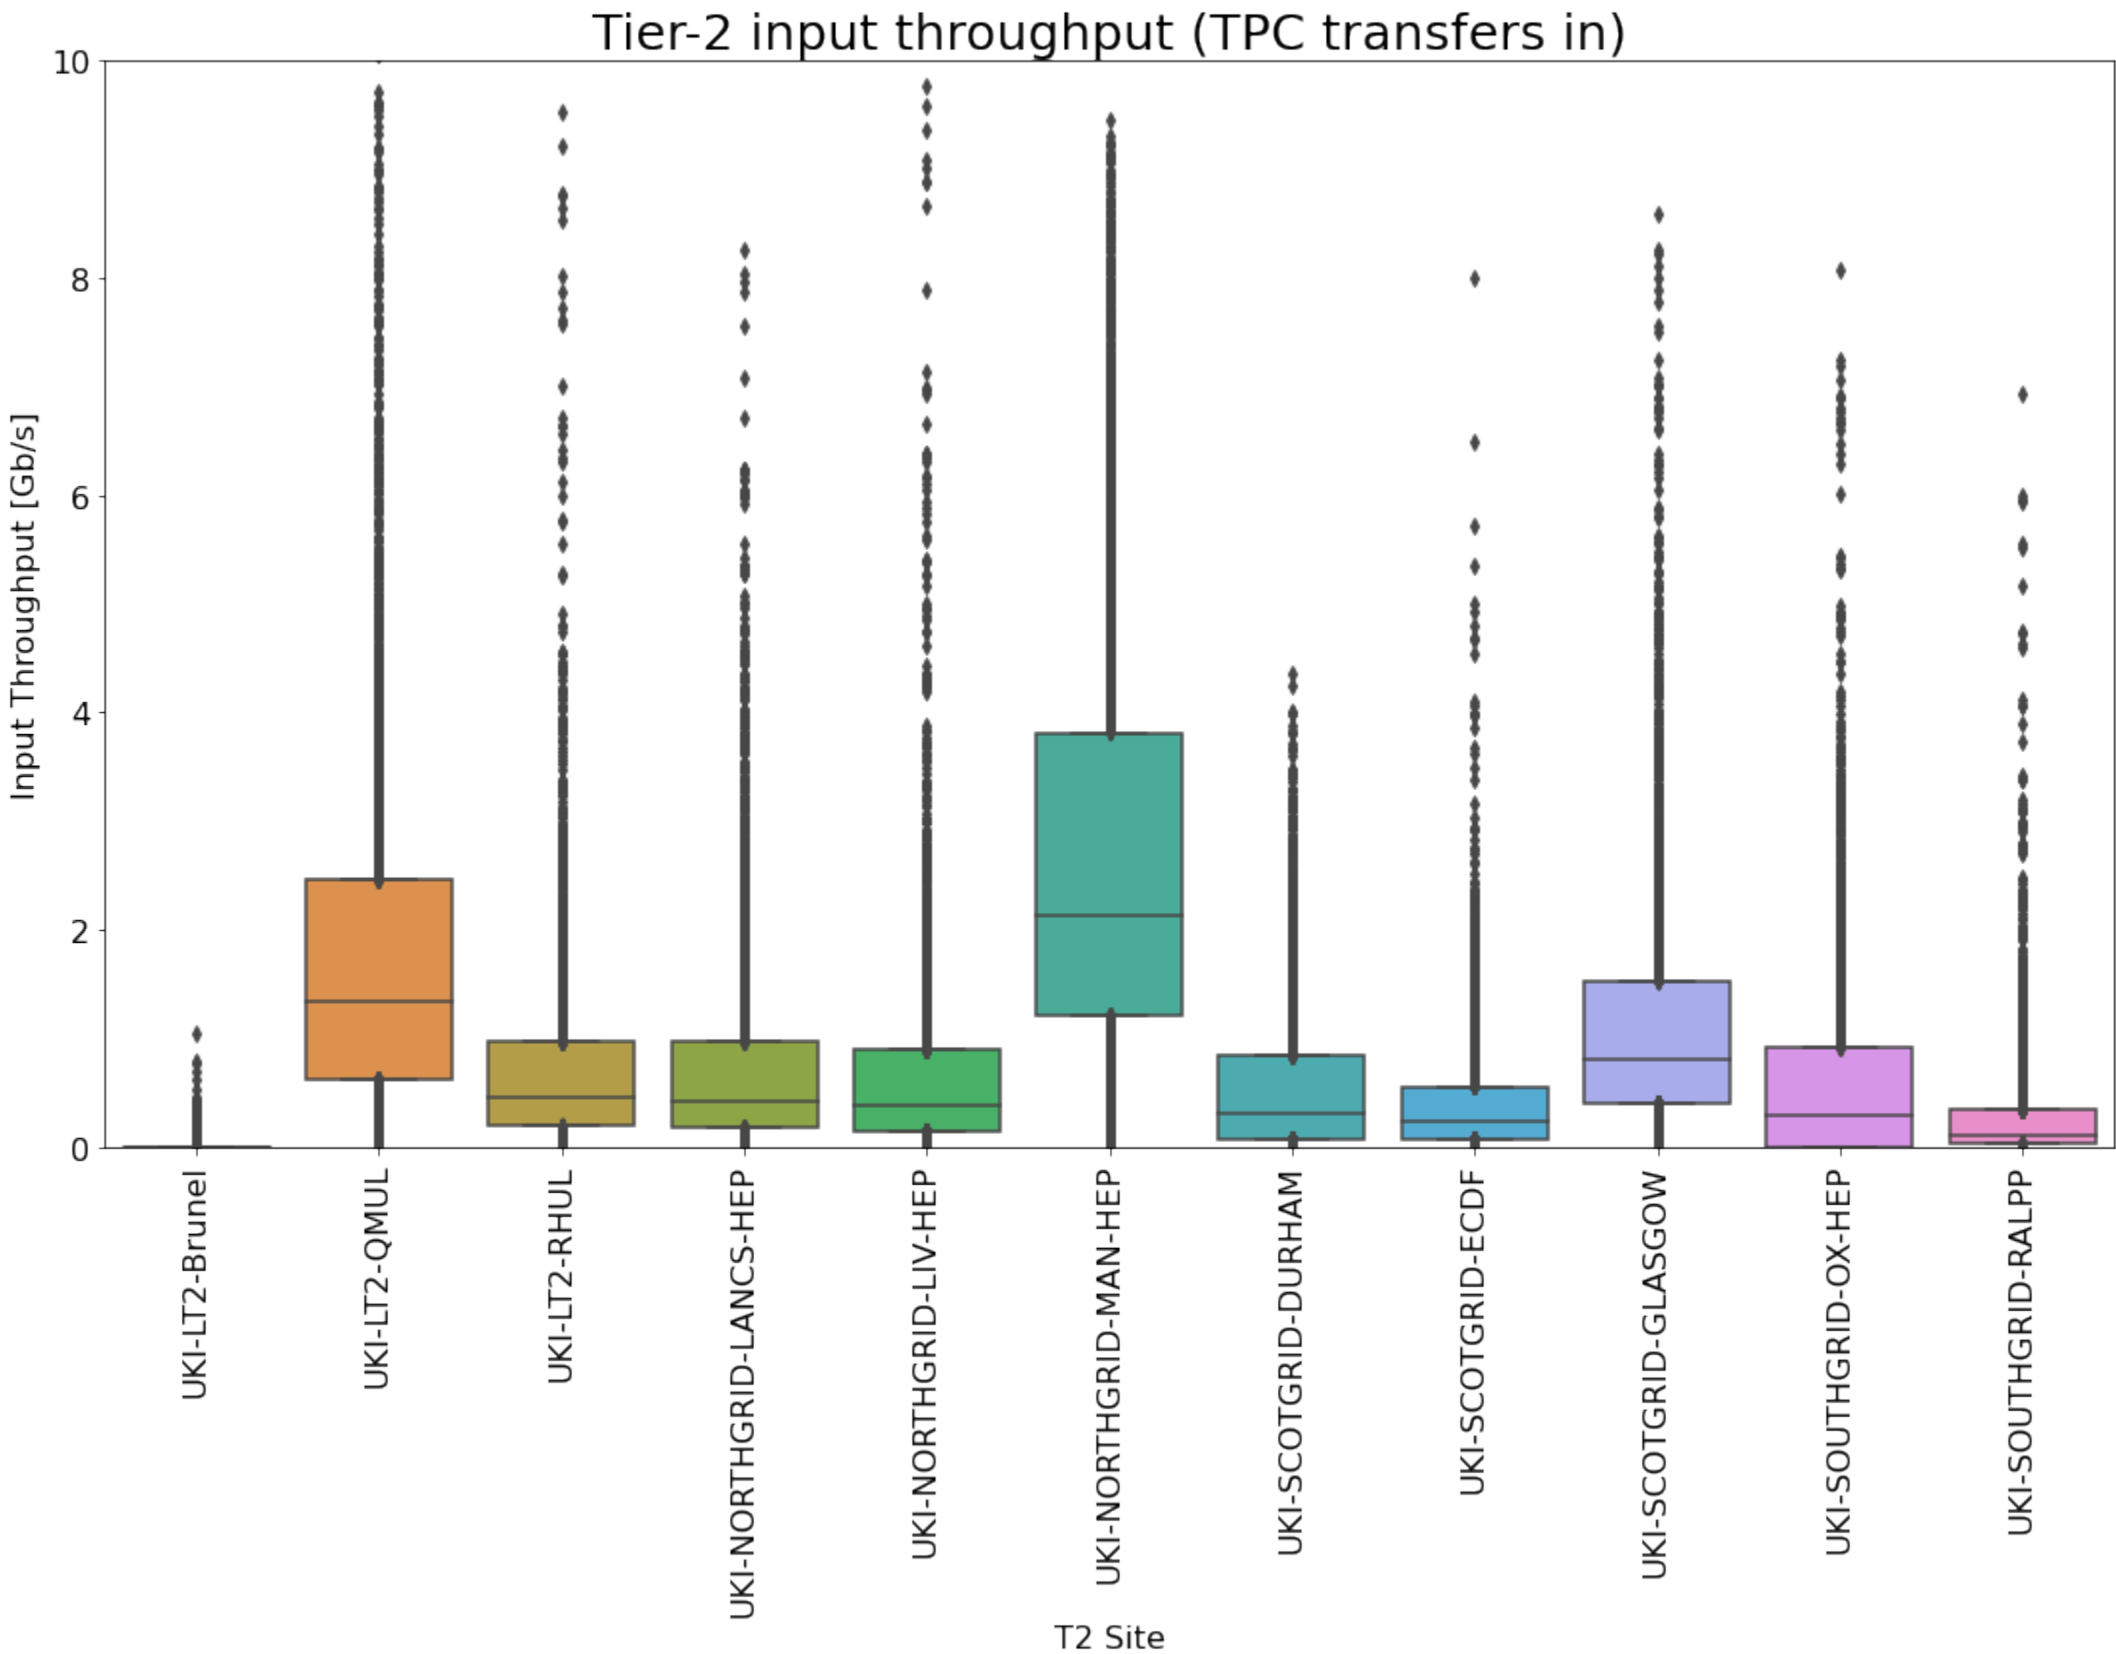

- Currently UK holds almost 30PB of data (in DATADISK),
  - Half contained at T2 sites
- ~30PB on TAPE at RAL
  - (17PB for MC, 13 PB for Data).
- During the last year, completion of number of major storage tasks:
- Glasgow:
  - Decommissioning of DPM storage (for ATLAS)
  - Completed Ceph commissioning; moved to production
- Oxford:
  - Migration to a storage-less site:
  - RAL operating as endpoint;
  - investigating XCache
- Note; storage decommissioning is  $O(3-6)$  month operation.





# Data throughput

- Data throughput into Tier-2 sites over last year (averaged over 1 hour intervals):
  - Bottom left plot: Transfers via FTS / TPC (e.g writing into DATADISK)
  - Bottom right plot: Including writing of data to local disk (or direct IO) for jobs
  - Right; Transfers from RAL to Oxford Xcache (max 9.6 Gb/s, top 25% of transfers around 4+ Gb/s.







# Moving to WebDav

- GridFTP transfers ~ deprecated;
  - WebDav protocol to be used for WAN and (ideally) LAN transfers
- dCache, DPM, Storm technologies all enabled
  - XrootD, main functionality provided with 5.2
- Echo / XrootD+Ceph functionally available, not yet deployed into production
  - Tuning of block / buffer sizes may provide throughput optimisation.


DOMA Stress / functional tests on RAL test gateway



- Fraction of transfers to / from UK Tier-2/3 sites by protocol

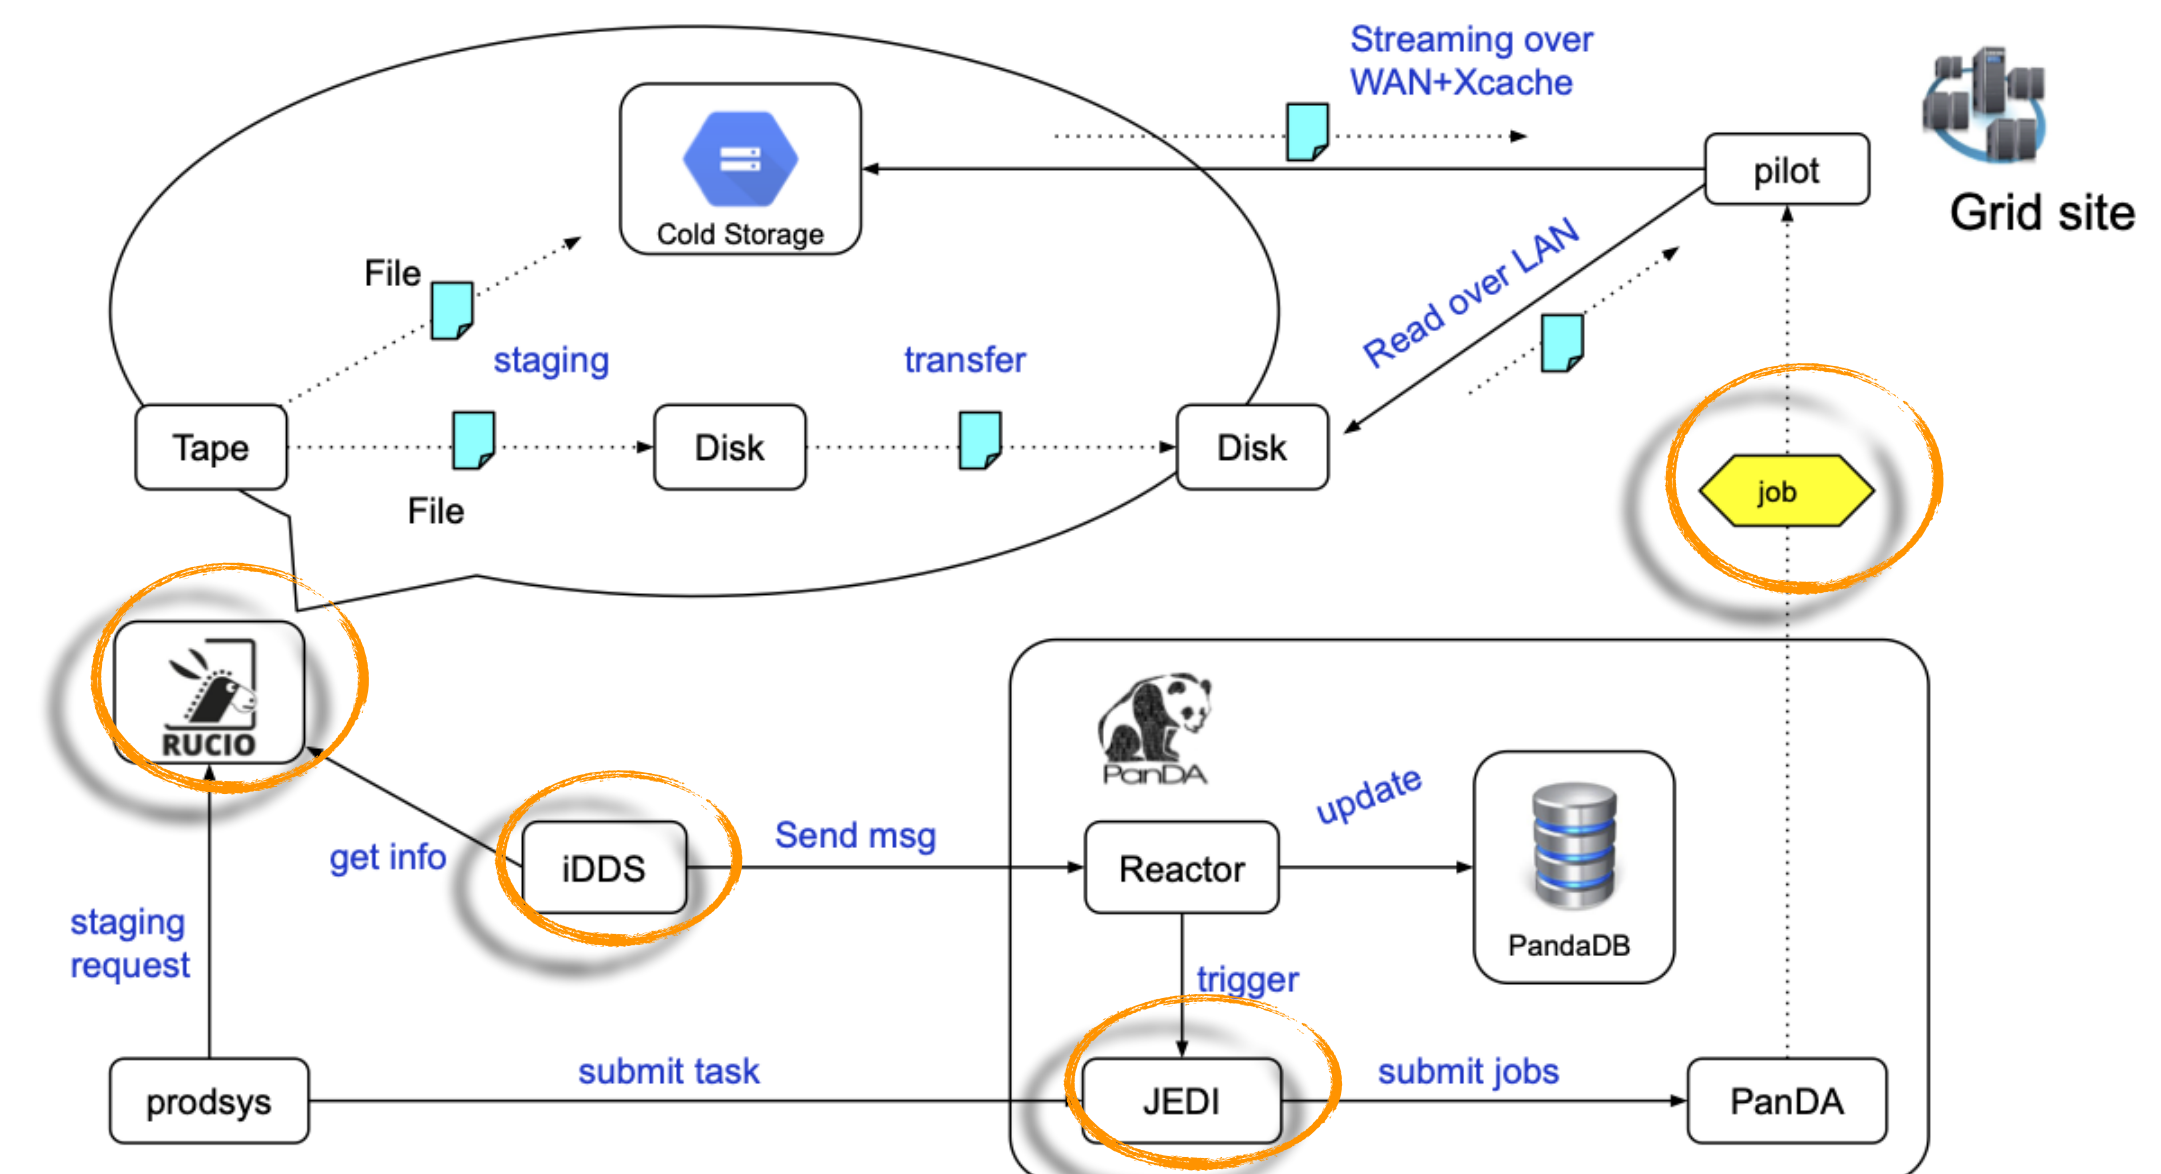



https download speed vs cache pagesize for RAL dev machine



# Data Carousel

- Data Carousel:
  - Provide orchestration between workflow management (ProdSys2 and PanDA), DDM (Rucio) and tape endpoints
  - Enable large (bulk) processing campaigns by staging a fraction (e.g sliding window) of data to disk at any one time.
- iDDS (Intelligent Data Delivery Service) introduced:  
Interacts between Rucio and JEDI  
(Job Execution and Definition Interface) to release tasks (jobs) for partially staged datasets
  - Decouples (asynchronously)
    - data pre-processing, delivery, and main processing
  - Also being exploited for Hyper-parameter tuning tasks and in non HEP fields.






- Data Carousel in use for some time now for production.
- ATLAS moving to larger file sizes stored on Tape, O(5–10)GB.



# Reconstruction Updates


- Athena reconstruction in release 22
  - new for run-3 and upcoming run-2 reprocessing
  - Move from MP (Multi-Processing) to (MT) Multi-Threaded model.
  - Plot (right) illustrating better sharing of memory from MT
- Throughput (plot bottom right), showing gains between r21 and r22:
  - improvements in track selection / optimisations contributing.
- On a ~ typically 8 core job with 16 GB allocation MT allows for better utilisation of resources:
  - e.g in r21 a 6-worker MP job just fits into the footprint.
  - 8 thread MT in 21 yields ~ 70% improvement in throughput.



# Simulation updates

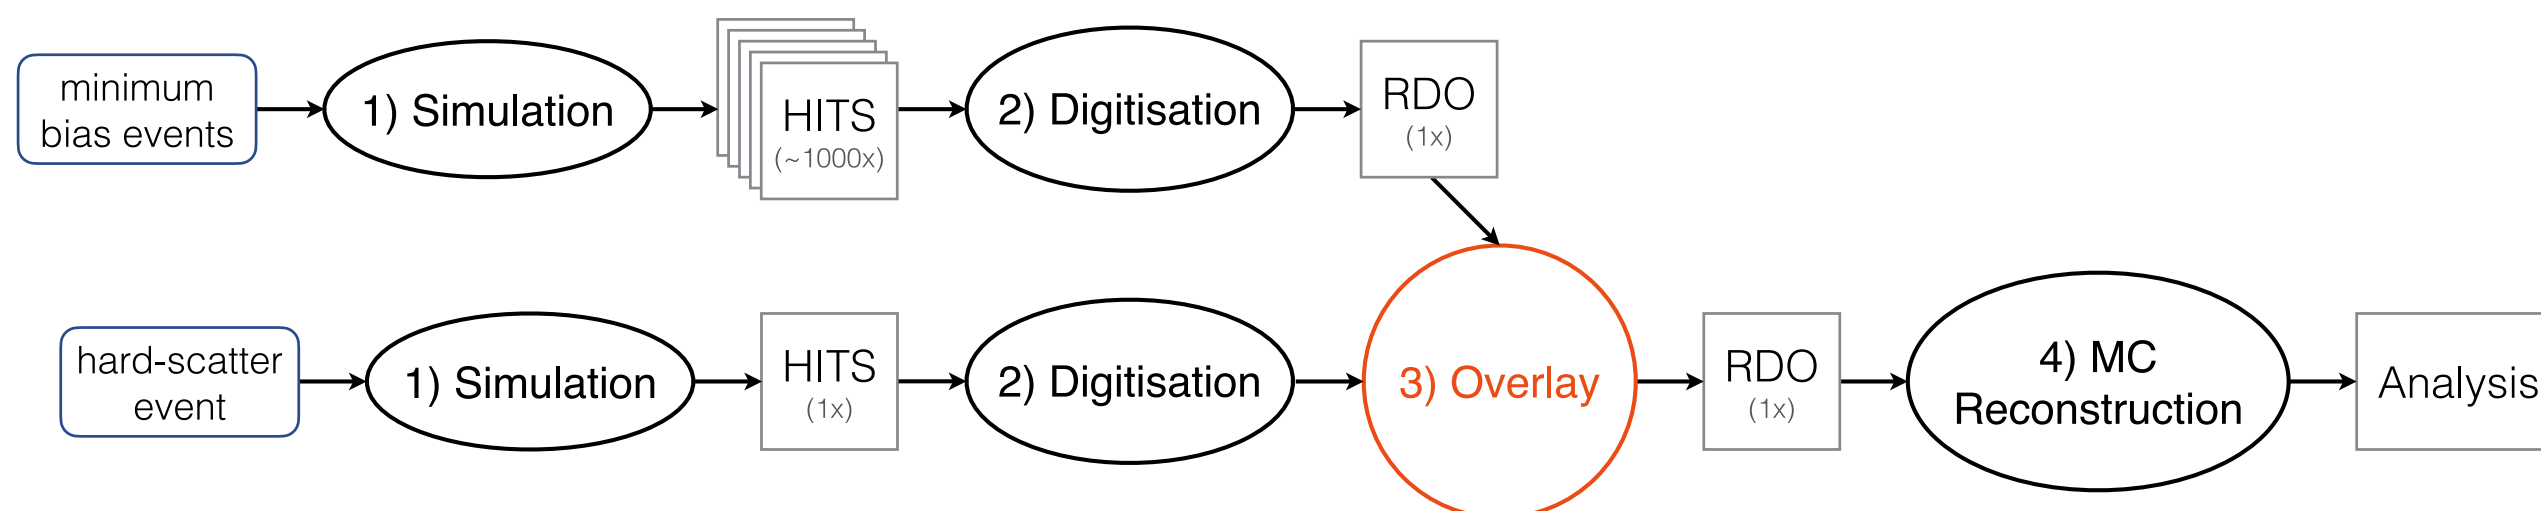
- Hard-scatter (process of interest) is superimposed on number of underlying ‘pile-up’ events in same collision process.
  - Up-to ~ 70 additional interactions during run-2 data-taking
  - (Even more complex as events in pre- and post-bunch-crossing may impact trigger response, etc.)

## Current simulation




- Min bias and hard-scatter combined at HITS level for digitisation
- MC+MC overlay approach:

- Large sample of already digitised mixed min. bias events overlaid on hard scatter event.



**arXiv:2102.09495**

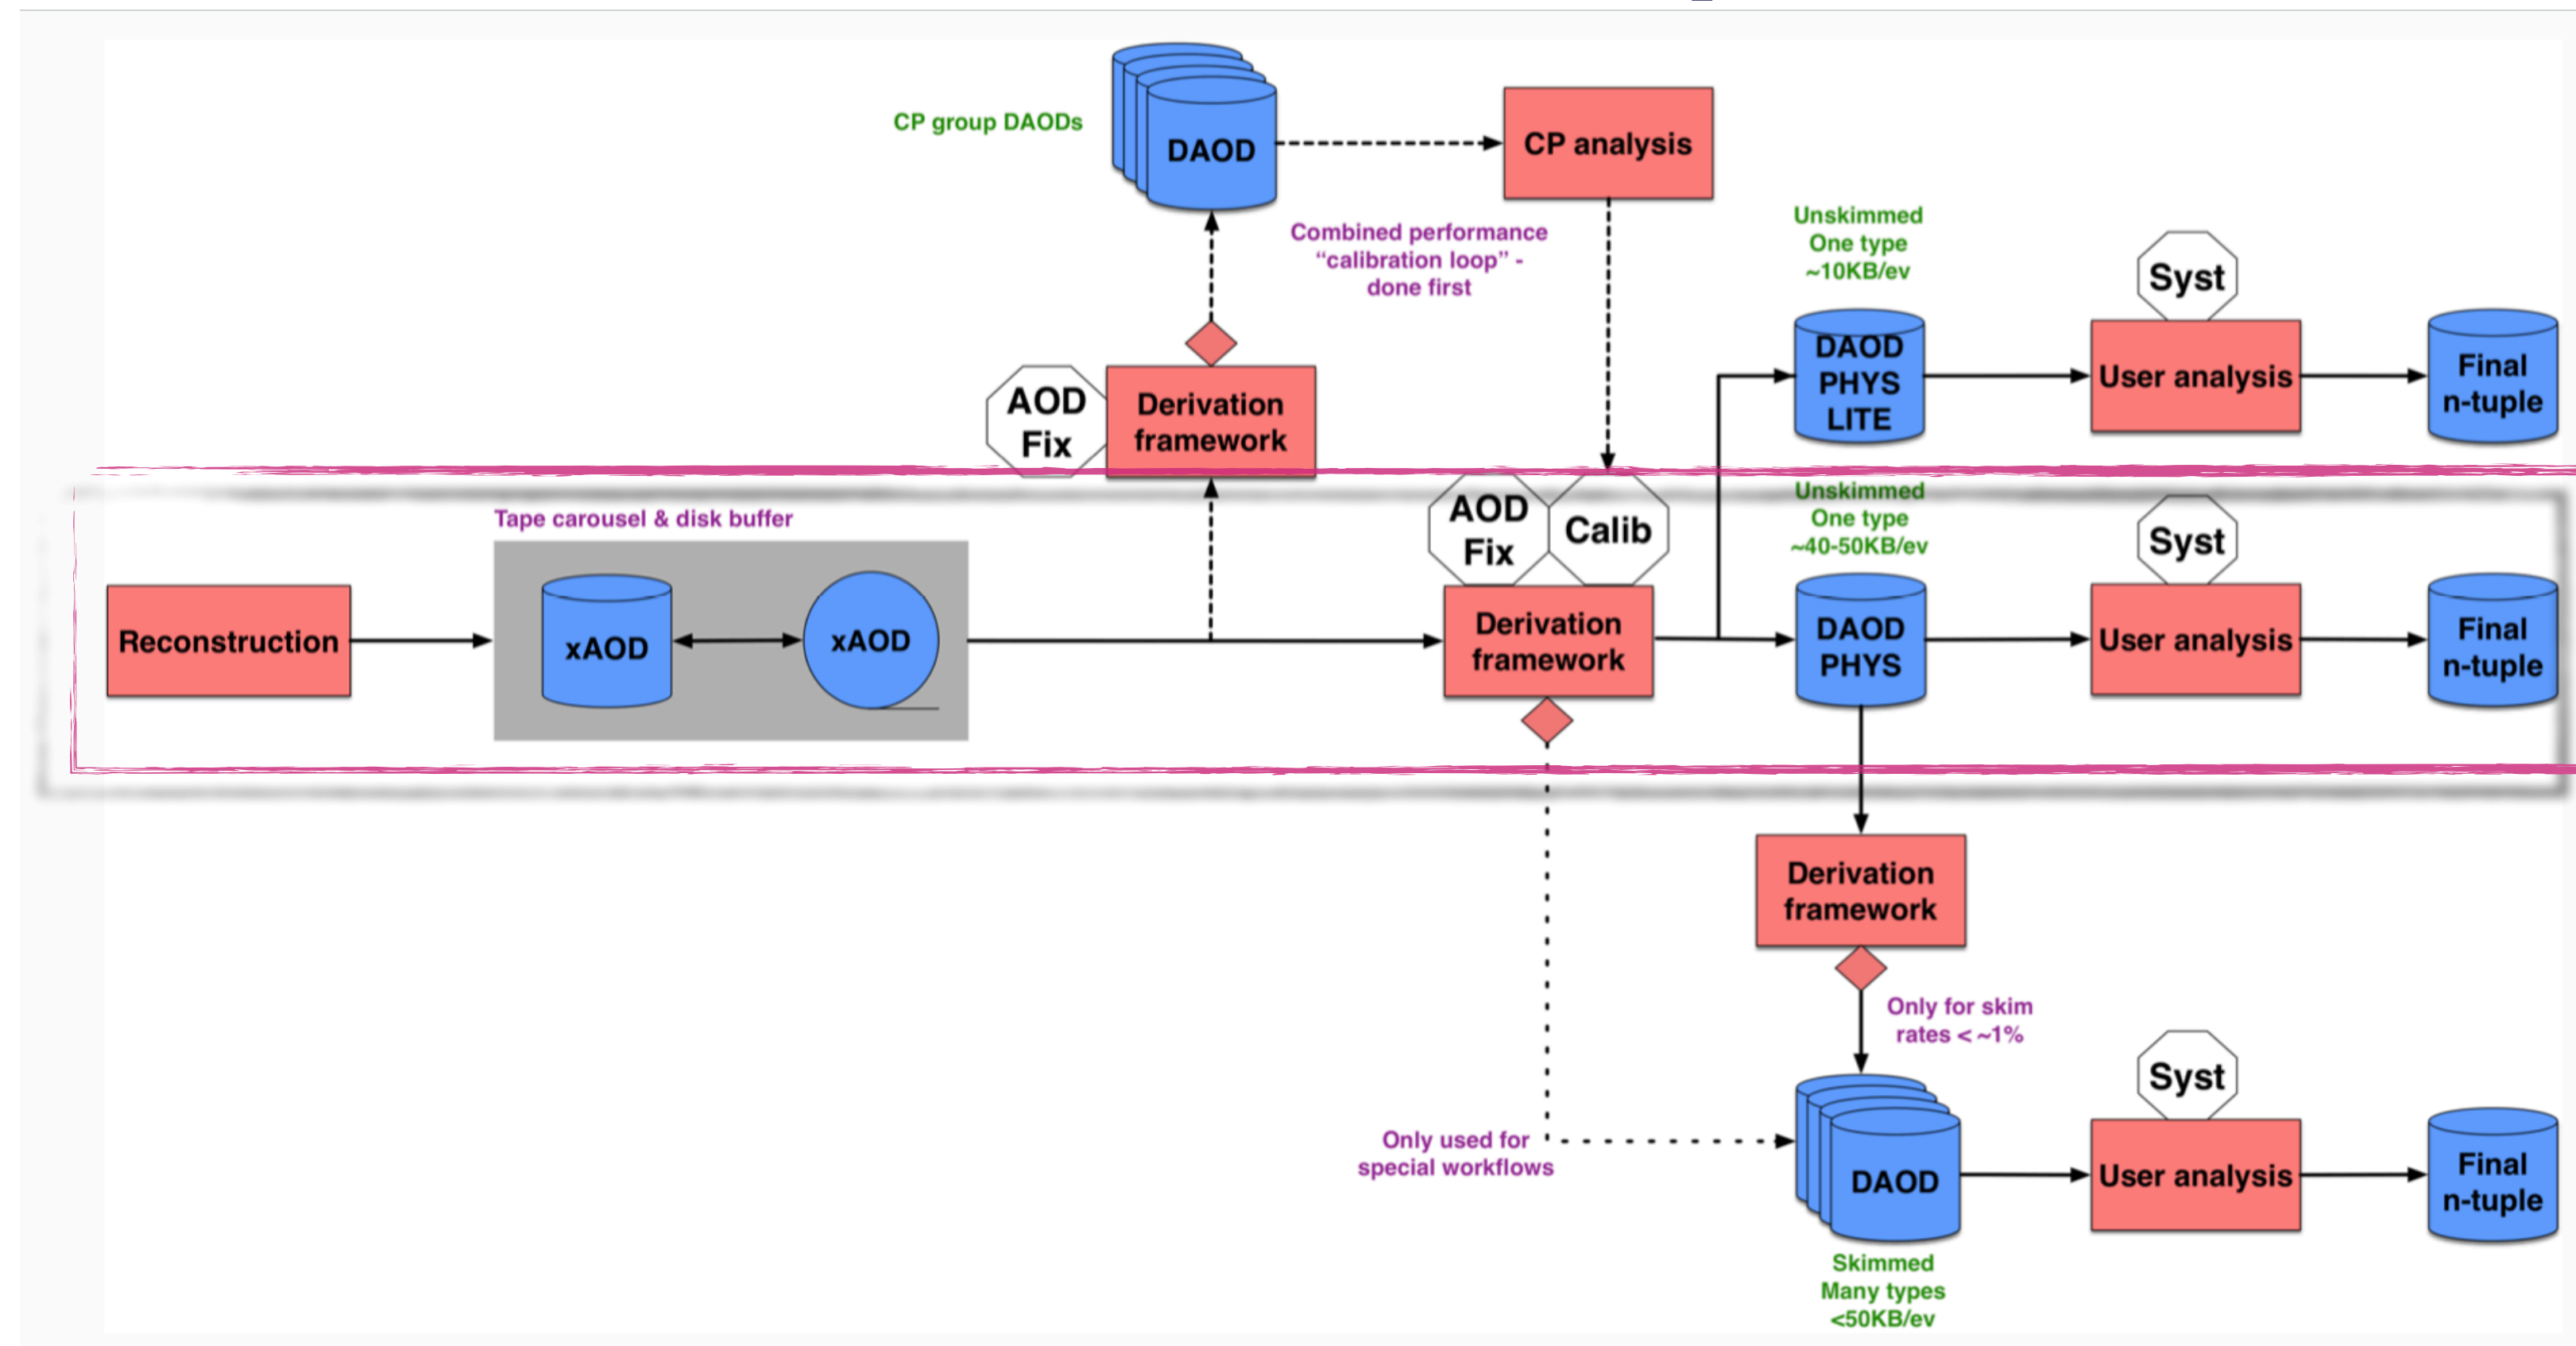


- Reduction of CPU by 20% anticipated
- RDO (minbias) step performed at sites capable of higher resource requirements,
- Overlay step somewhat simpler on resources: Smaller Minbias files and with sequential read access;
  - Can be performed on larger set of sites, with prestaging of data



# Derivation format and Workflow updates

- Primary analysis format (AOD) too large for general analysis:
- Run-2;  $O(100)$  formats of derived AODs of  $O(1\%)$  initial volume.
- For Run-3 :
  - Aim for single format, appropriate for most analyses, and prototype run-4 super-condensed format.
  - Unskimmed, but some reduction of information (e.g removed tracks below given pT).
  - Lossy compression potentially available for truncation of floating point bits in certain observables.
- DAOD\_PHYS:
  - $\sim 50\text{kB/evt}$
- DAOD\_PHYSLITE:
  - $\sim 15\text{kB/evt}$
  - Containing calibrated objects only
  - Potential for better hit rates in Xcaches, or whole years at single sites.
- Current DAOD formats
  - May exist for a while, while analyses complete
  - Some specialised DAOD formats will remain, for CP groups, and physics groups not able to utilise standard selections.
- AODs:
  - Available primarily on TAPE, with coordinated campaigns for new DAOD productions



Format	Event Size [kB]	Nominal total size / collision year / version
AOD	600	
DAOD_XYZ	40-450	
DAOD_PHYS	30-50	$O(1.3)$ PB
DAOD_PHYSLITE	10-15	$O(0.5)$ PB

# Summary

- Run-3 rapidly approaching;
  - Number of Data challenge-type tests to begin (See talk from Alessandra Forti)
  - Full Run-2 reprocessing shortly to get underway
  - MC generation campaigns for run-3 to start ~ End-of-year.
- Transition to tokens;
  - VOMS deprecation during run-3
  - Token-capable services should generally be in place for start of run-3
  - Discussions are still active on tokens management and workflows
- Heterogeneous computing; ATLAS exploring and utilising various workflows (eg. Commercial Cloud, HPC, GPUs for hyper-parameter optimisations, etc.)
- New data formats for reduced analysis data-size footprint for run-3 (and run-4 prep. developed).
- Updates to reconstruction software and MC min-bias overlay; reduction CPU, IO consumption
- Feasibility of exploiting ‘unreliable’ storage and QoS under study (See talk from Rob Currie)
- Increased usage of TAPE for AOD storage (with Data carousel model)
  - Disk space will remain at a premium going into HL-LHC

