


Dilepton $t\bar{t}$ candidate event







Science and
Technology
Facilities Council

GridPP48: ATLAS

James Walder
2 September 2022

ATLAS Status


- Already surpassed 10fb^{-1} integrated luminosity;



Tuesday 23 Aug evening the cooling tower in P4 (SF4) failed
Experience from March event is a total downtime of 4 weeks

- Warm-up 10 days
- Cool-down 4 days
- Conditioning 2 weeks


- Recording efficiency improving and following trajectory from 2018



ATLAS UK Overview: Compute


- Average compute ~ 610k HS06 over last 6 months
- As usual, T2 federations consistently delivering over their pledged values.

	max	avg	'22 Pledge
UK-T1-RAL	359 K	185 K	192k
UK-NorthGrid	229 K	169 K	74.2k
UK-ScotGrid	231 K	149 K	50.0k
UK-London-Tier2	173 K	78.2 K	34.7k
UK-SouthGrid	41.0 K	29.3 K	18.9k



- A few periods where one site in a federation has suffered issues;
 - Other sites maintained the pledge for the T2 fed.
- Next milestones:
 - ATLAS UK to reach 1MHS06 ...
 - Aside from new deployments (e.g. RAL),
 - this is a zero-sum game, and currently requires some other VO to suffer misfortune.
 - However - getting ever closer to being a genuine milestone.


Currently* ~ 961k max HS06 (hourly average)




UK composition of jobs

- UK (~ broadly similar across other Clouds (Grid sites))
 - In terms of jobs submitted:
 - ~ One third jobs from User Analysis
 - (~ half are User + group Analysis)
 - MC Evgen ~ 20%
 - For used wall clock:
 - User (+group) analysis ~ 25%
 - Evgen ~ 30%
 - MC simulation and reco accounting for majority of remaining resources in this period (last 6 months)

Numbers of Jobs



Jobs by Wallclock

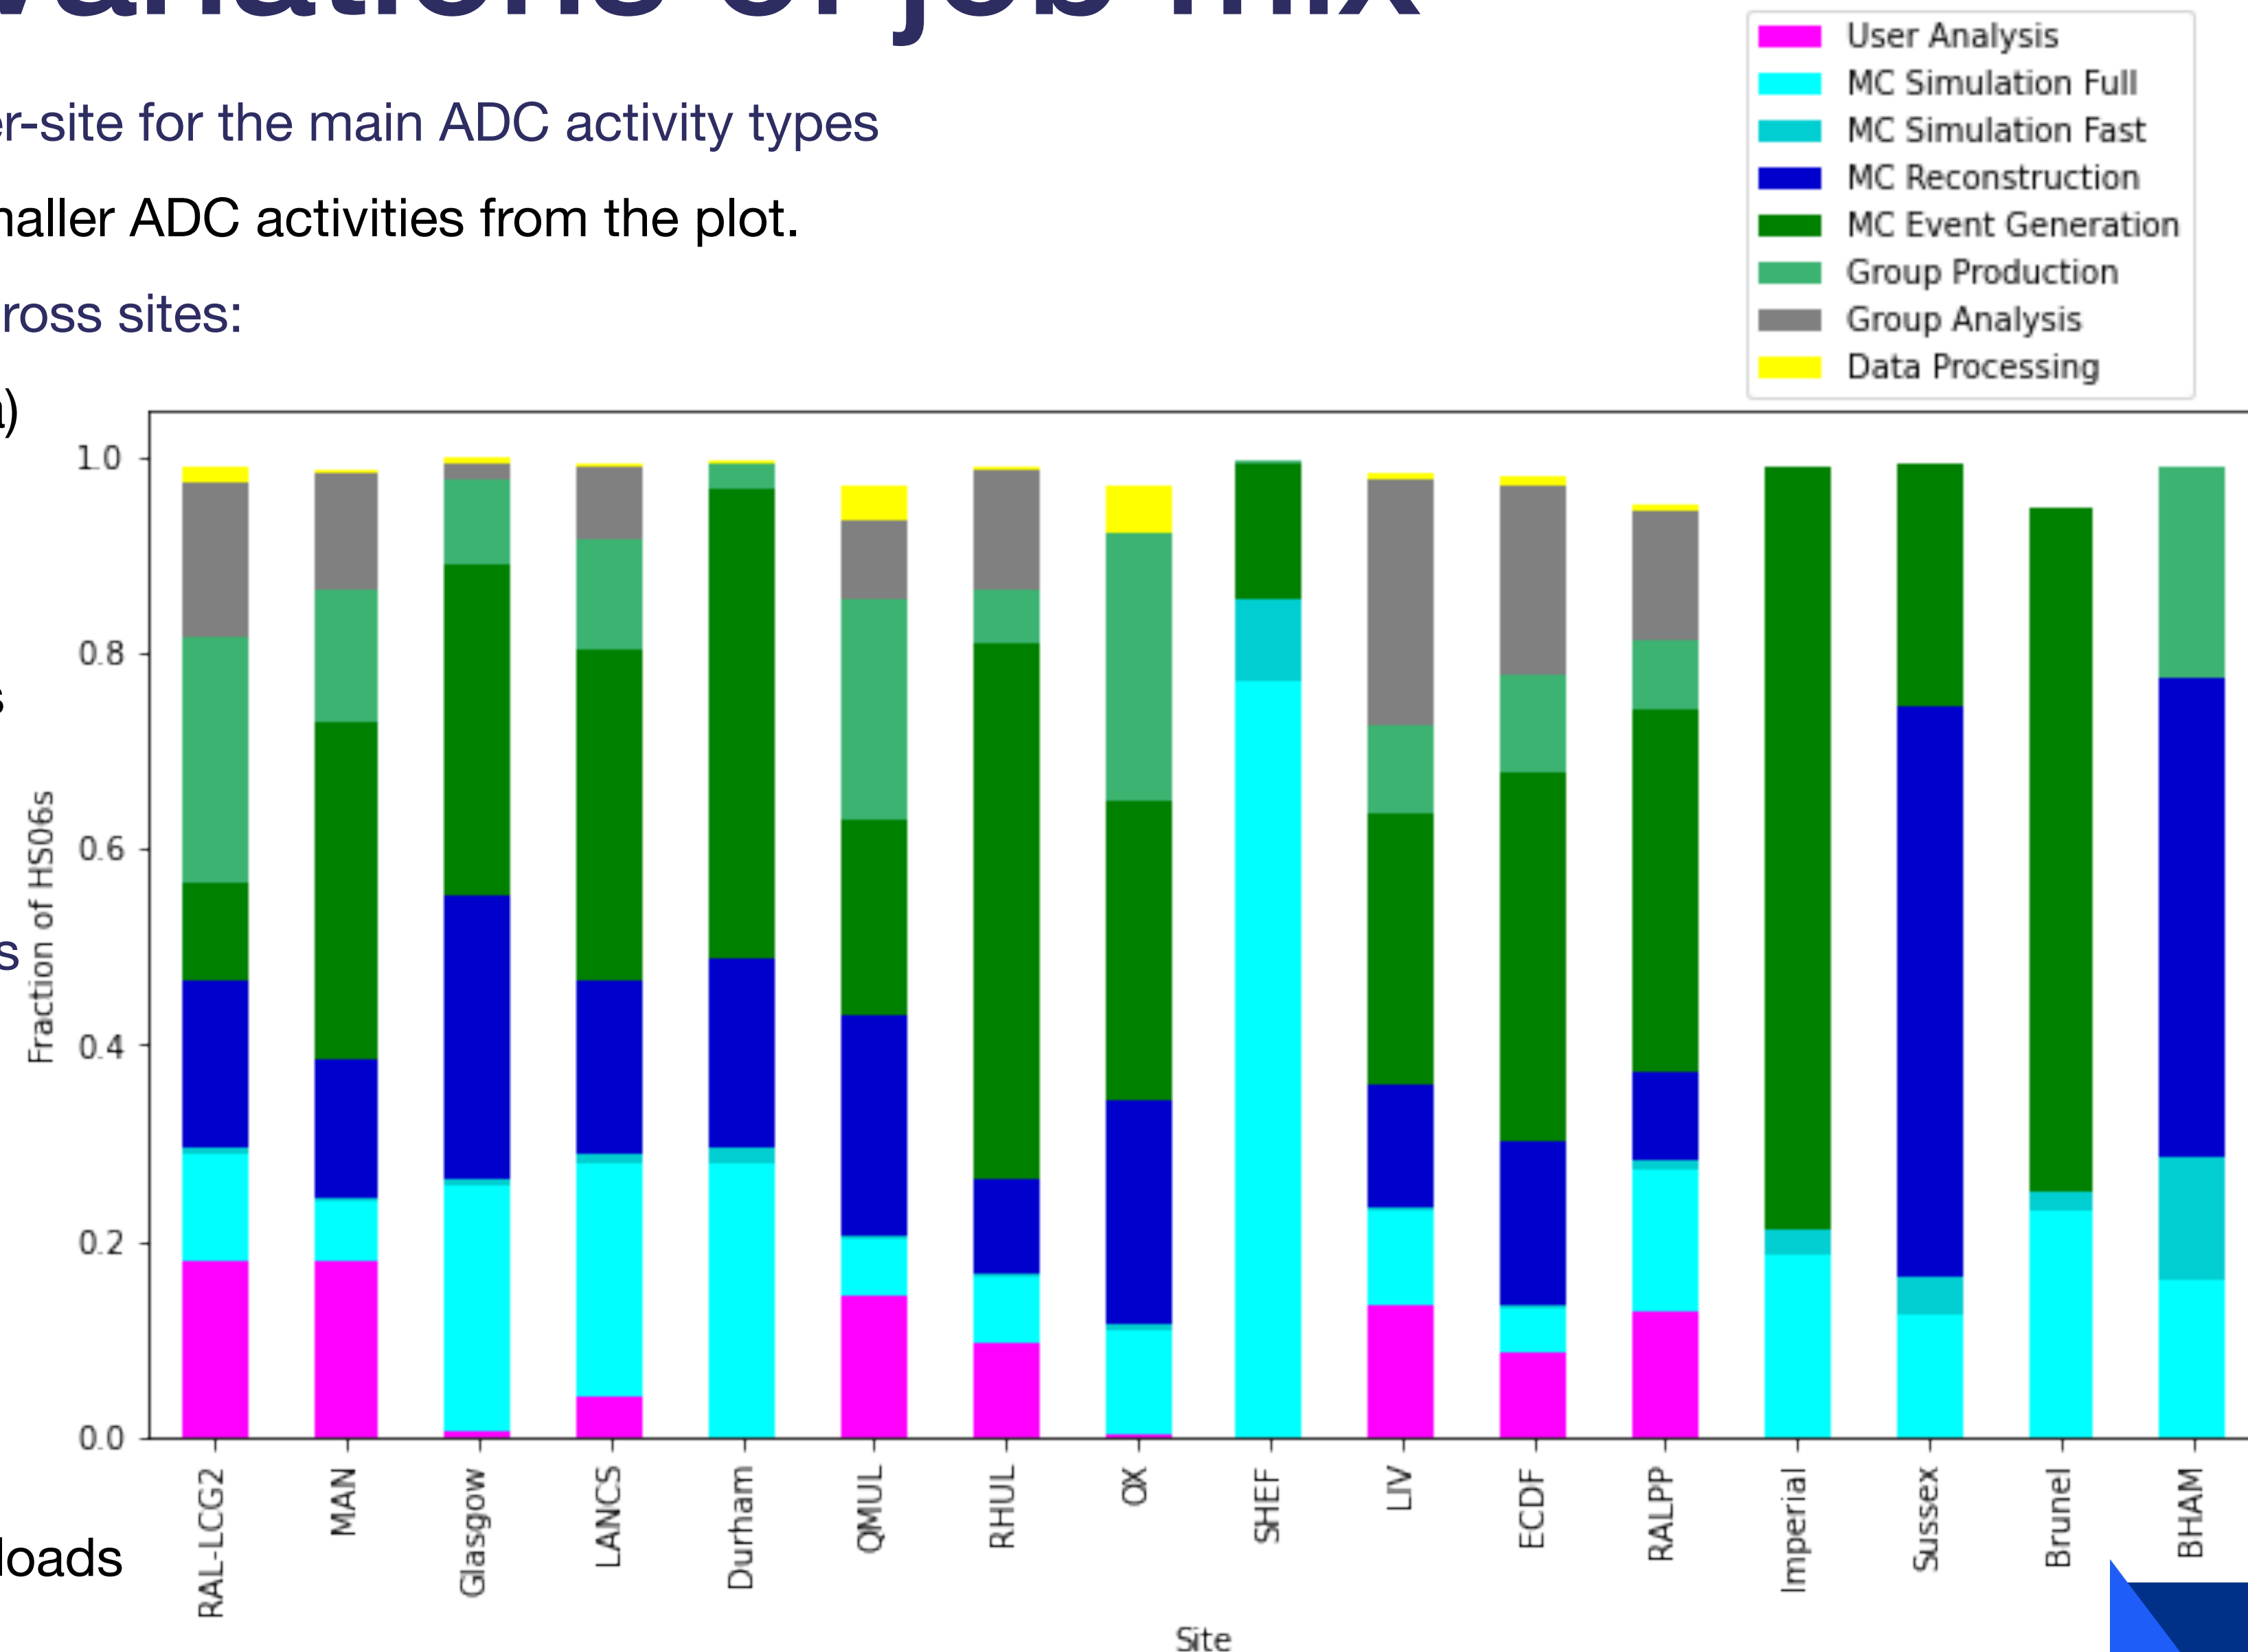


Site variations of job mix

- Plot of the fraction of HS06s resources per-site for the main ADC activity types
 - Normalised per-site, and dropping smaller ADC activities from the plot.

- Several potential reasons for variations across sites:



- Data locality (i.e. already located data)
- Unified / Prod / Analysis settings
- CRIC overrides: e.g.
 - Restricting less IO intensive tasks to certain sites, either by:
 - Explicit job type
 - Expected task io requirements
- Brokering requirements:
 - Memory, IO, timelimit
 - Available slots of single / score
- #TODO: review the current explicit (and hopefully implicit) settings for job type to UK site mapping
 - Ensure allocations of (intensive) workloads are still 'fairly' distributed.




Data Transfers

- Source: [FTS status board](#)
 - T1: Usage is a volatile function of demand
- Intra-UK: 3.7 PB of FTS transfers (6 months)
 - (Excluding intra-site transfers (e.g. multihop))

Volume [TB] transferred (last 6 months): Intra-UK






Volume [TB] transferred (last 6 months): T1/0-T1/0



- 6 months
 - ~3PB from Cern to RAL
 - ~2PB from RAL to CERN




Site Datadisk Occupation Rates (1)

- A few sites recently provisioned new storage:
 - Either as a new SE (e.g. Lancs)
 - as an addition to existing (QMUL), or after a large (atlas) deletion operation (e.g. RAL)
- New storage has a 'commissioning period', then appears to settle:
 - 5PB/m for RAL-LCG2,
 - 1PB/m for QMUL (~0.7 PB/m for second increase).
 - 0.7PB/m for LANCS-CEPH



Site Data Occupation Rates (1)

- LANCs:
 - Data input for production jobs, data consolidations and outputs from jobs
- QMUL:
 - More varied inputs: Data for Analysis, consolidation and “nucleus” collection of data
- RAL:
 - Significant “staging” (e.g. multihop), otherwise mainly analysis and production inputs.
- No clear conclusions; but interesting to understand better the correlations between rate, status, compute, bandwidth and other (?) factors.



Colours not consistent across the plots


RAL Tape Usage

- ~ 15PB active data in DATATAPE
- ~ 23PB in MCTAPE
 - 38 PB of used storage
- 40.289 PB for 2022 Pledge
- Expect to need to go into the 2023 pledge early and a request is in preparation.
- If meaningful, run another deletion campaign.

Archiving to RAL

6hr binning




Transfer Throughput



Staging from RAL


6hr binning

Transfer Throughput



Improvements to Multihop: FTS 3.12

- Multihop archive to RAL-LCG2 Tape generally working well.
- As overwrites to TAPE are not allowed / enabled, previous FTS versions would also deny overwriting files in any of the intermediate hops
 - Transfer failures into Echo (first hop), would then result in files that could not be removed by usual means.
- FTS 3.12 brought in a number of changes:
 - Python 3 only bindings
 - Flag in the FTS configuration / db for CTA-based tape sites:
 - File eviction from tape no longer relies on “cern_tape_archive” string in the Site name (but relies on FTS instance to know that Antares is a CTA endpoint)
 - And, a new flag to allow to overwrite intermediate hops (which rucio understands)
- Had previously required regular manual cleanup to remove any files the failed in the intermediate hop.
 - Clear change once fix in place




‘POSIX’ Read Access for Batch Jobs

- Current model of data access typically proceeds via a gateway:
 - Batch jobs reads (copies) and writes from to the storage endpoint (e.g. XRootD gateway)
 - Sites with storage (posix-like) mounted spaces on the worker nodes should be able to make use of this approach:
 - Reduce load / latency via the Endpoint/Gateway
 - Direct access already available for Storm (e.g. QMUL), but needs to access additional metadata.
- Motivation particularly with cephFS + XRootD setups (i.e Lancaster)
- Datadisk is mounted read-only (and with a read-only (Apptainer) bind-mount in CRIC).
 - Reads go via the mount (and not through the Gateway.
 - Writes (for auditing, etc ...) are still performed through the XRootD endpoint.
- Instead of “rucio get” downloading a file, a symlink to the file is created instead.
- Tested and implemented in rucio 1.29.X as a posix.Symlink implementation to the posix protocol (file scheme)
 - For ATLAS, awaiting 1.29 to be default within the job, and an update in the pilot to allow “file” protocol access.

High IO Evgen (1)


- Recent incident of excellent inter-operation between site and Experiment.
 - RAL started to observe poor CPU efficiency and failing jobs.
 - Mostly affecting older HDD-based nodes, new nodes (with SSD) also saw some IO-wait.
 - Tom Birkett identified ATLAS jobs, and identified that a file was constantly being (re)-written with 40B of data (~ 3kHz).
 - Worked out that Sherpa EVGEN jobs were creating high IO load:
 - (Storing random numbers in this 'temporary file').



- Worked on fixed (using the shared memory tmpfs space /dev/shm) which then implemented and deployed to production by MC experts.

High IO Evgen (2)

- (Left) pre-fix, high volatility across most sites
- (Right) post-fix, less volatility and higher median values



- Latest campaign of MC generation moving to non-CLHEP RNG method of storing random numbers:
 - (RAM-based storage)



Data placement and caching (1)

- A few methods of testing data placement underway in UK:
 - Caches, Volatiledisk, virtual placement
- **Caches**, e.g XrootD XCache:
 - Either transparent (i.e. hidden to ATLAS) within a site, or exposed and registered in CRIC
 - Oxford XCache connected to RAL-LCG SE.
 - Remains a single-point of failure, but generally works “well”.
 - Data workflows suggests low expected cache hit rate, so most potential is in latency hiding (?).
 - To test with more recent hardware shortly.
 - May want to test with (artificially) harsher environments (added latency / packet loss / network saturation effects) ?
- **Volatile** Disk:
 - Data registered in rucio, but potentially with different QoS to ‘standard’ Datadisk.
 - ECDF participating with test storage as part of QoS efforts.
 - Manually degraded some fraction of files
 - Not aware of much recent activity in this area, but remains an interesting problem..

Data placement and caching (2)

- **Virtual Placement:**

- Xcache-based rucio-aware solution to
 - Rucio can query for known replicas, on job assignment to a (VP) site, rucio can provide a (cache-prefixed) URL to the data.
- Some US sites (and Prague) previously tested
- Birmingham participating;
 - Work from Mark S (together with Ilija Vukotic) to avoid dependency on Slate
 - Ran ok, but did not fill all potential slots
- More recent work from Bham. in trialing their non-VAC system
 - Working well (with a few final 'tweaks' to investigate).



- Trying to enable VP back at BHAM again, but in non-VAC queue.
 - In current configuration, did not appear to be able to fill a site with jobs, but currently could be useful to supplement, and may be more effective with further improvements in config.
- VP continues as an active development, with an aim to be included (fully) within rucio.
- Interest in ATLAS also in trying to see if data (pre-)placement can be improved.

Token support within ATLAS (1)

- Run-3 -> HL-LHC to see the deprecation of VOMS in favour of tokens
- ATLAS timelines (as for other WLCG experiments) driven by technological advancements and maturity of infrastructure, along with site ability to support these (fundamental) changes (and of course LHC timeline).
 - Macaroon-based (storage provided) tokens already used for some time for FTS transfers
- GridFTP ~ now removed from all (storage) infrastructure
- Main points from WLCG timeline:

M.6 (Mar 2023): Some storage endpoints provide support for tokens (at least one per service type).

M.9 (Mar 2025): Grid jobs use tokens for reading and stageout.

M.10 (Mar 2026): Users no longer need X509 certificates.

- To enable the testing of many of the additional components needed in the WFMS (next slide), need reasonable storage support for tokens (e.g. latter 2023).



Token support within ATLAS (2)

- WLCG compliance testbed used to verify / confirm storage implementations.
 - XrootD 5.5.0 brings necessary update to storage.create (which should make it ‘fully compliant’ with current setup).
 - (RAL participating in this, but need dev machine back from other work to confirm all tests are now passing).
- Near-term plans / updates:
 - HTCondor end of GSI support => token submission to HTCondor-CE & X.509 submission via ARC-CE REST
 - Transition from VOMS legacy servers to IAM VOMS: clarification of support level for IAM VOMS:
 - IAM VOMS needs to be treated as a critical service.
 - WLCG Data Challenges 2024 with tokens.
- With established token support in storage, can add more WFMS related items (~2023), e.g.:
 - designing token workflow for ATLAS jobs (harvester & aCT+ARC-CE data staging)
 - Expect discussion of related topics during <https://indico.nikhef.nl/event/3612/> (ARC/HTCondor CE Hackathon).
 - Working ‘rucio upload’ and ‘rucio download’ to be able to test real jobs with tokens
 - Support for {read,write}_lan with tokens (for xroot protocol),
 - direct-io comes with additional challenges
 - PanDA already supports tokens; but testing needed once full workflow chain is available
- A number of other non-WMFS services also need token implementations; e.g.:
 - accessing AMI with tokens, accessing CRIC with tokens, ...
- The User Interface for (/experience with ?) tokens; hoped to be discussed in future WLCG AuthZ WG meetings

Tokens: Implications for sites

- Campaigns:
 - September 2022:
 - Necessary updates of HTCondor-CE with tokens & ARC-CE with REST
 - [GGUS: WLCG Token support campaign](#)
 - HTCondorCE, Arc-Rest tests:
 - <http://novastore.farm.particle.cz/ce/arc-rest/> , <http://novastore.farm.particle.cz/ce/condor-ce/>
 - hepgrid6.ph.liv.ac.uk and lcgce1.shef.ac.uk currently showing as 'red' on these tests.
 - **DPM EoL** – Likely to have Sites contacted under a “migration campaign” (GGUS);
 - expect the “dCache migration script” to be the de facto ‘recommended’ approach.
 - A driver for the future of Storageless sites in UK?
 - “Later” (~2023):
 - Storage upgrade campaign:
 - Sites to provide good token support
 - Sites to continue to support x509 VOMS
 - Paths accessed via tokens expecting to have a common namespace across sites; ‘might’ (in some cases) end up with different LFNs for path with tokens vs path with x509 (needs to be confirmed if this is a real use-case however).

Tokens: Open questions ...

- A few comments / open questions:
 - Tokens now being realised from a number of communities,
 - Some with less / no 'history' from x509
 - ATLAS (and WLCG) generally focusing on own use-cases for migration:
 - e.g. multi-VO rucio support via tokens?
- 'In the wild' usage of sites supporting multiple WLCG JWT, EGI AARC, ... token profiles?
- Continuing development of core functionalities:
 - e.g Submission of EGI jobs with EGI CheckIn tokens to the HTCondor-CE (not currently available)
 - HTCondor-CE with GSI support is EoL next March.
- Rucio (+new infrastructure) will orchestrate many of the complications around tokens:
 - Token exchange, refresh tokens, vaults, ...
 - Technologies may be demonstrated, but many *final* implementation / policy details remain actively discussed.
- Much new functionality is needed / being provided in the move to tokens; some legacy components will disappear (and some will need to coexist ...).

Summary

- Run-3 has started!
- UK continues to perform well overall, delivering (and exceeding) its commitments.
 - Operationally active times, with many upgrades (network, software, hardware, Data Centres, in progress / to be done)
- Storage remains a critical resource
- Usual Reminder: Storage decommissioning / migration has a long lead time:
 - Let us know in (lots of!) time; also applies to DC major works ...
- DPM End of Life is coming; UK is unlikely (?) to follow the suggested dCache migration route:
 - Lancaster XRootD+CephFS so far looks good; stress-test with mini data challenge ?
 - Aim of rucio posix.SymLink to facilitate read_Ian activities,
 - XRootD Redirector managing remaining load.
- (WLCG) Token support: ... lot's still do to / be understood and implemented and tested
- Preliminary ATLAS resource requests just submitted to the C-RSG for the 2024–2025 pledge periods
- Additional 0.5 FTE starting next week on GridPP (Jyoti - RAL PPD) in a Liaison role and will be getting up to speed with ATLAS (e.g. starting with cloud support efforts).

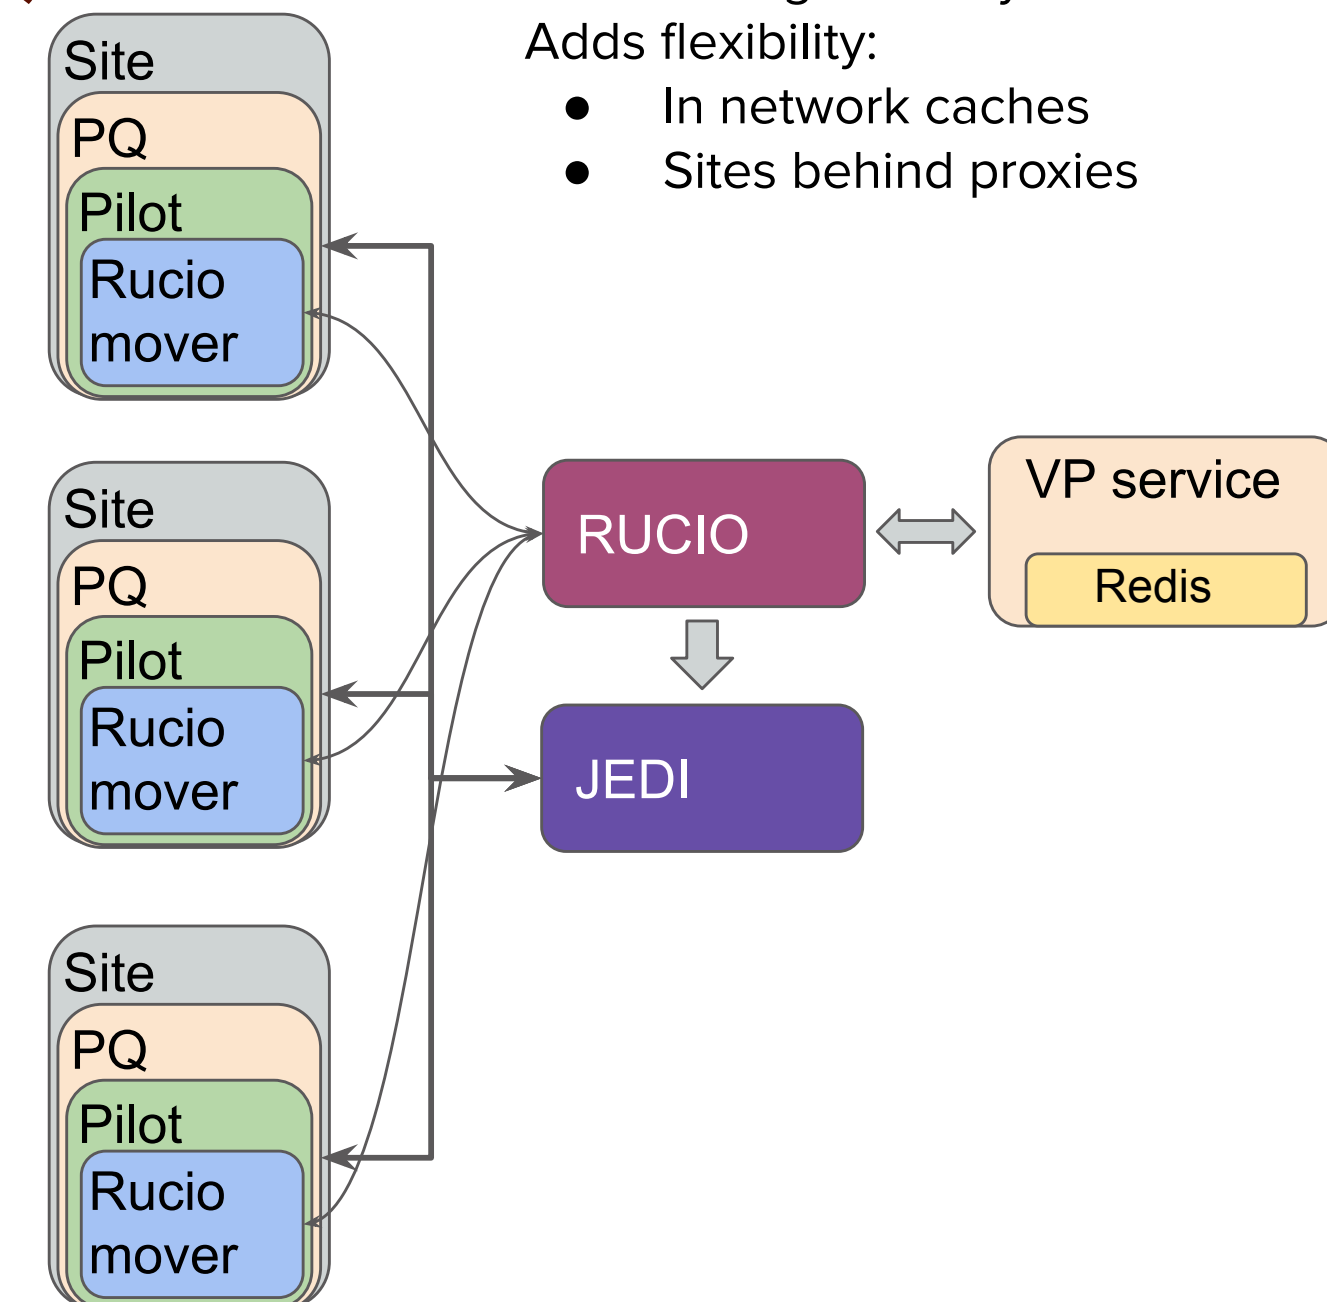


VP - Virtual Placement

1. The first time Panda needs to process a dataset, it gets assigned to N VP DDM endpoints in the same region.
2. Assignment is done randomly and each VP DDM endpoint probability to get the dataset is proportional to fraction of CPUs that its site contributes to ATLAS.
3. Datasets are not actually copied at any of these N endpoints but only exist in the “lake”.
4. Panda would assign job that needs as an input this dataset to the first site from these 3. In case site is in outage it would get assigned to the second site from the list. Once job is there it would access the data through the cache.

VP in operation (new)

- XCaches send heartbeats to VP service which keeps track of active caches and their sizes
- Rucio once per minute retrieves info on xcaches at each site.
- Rucio configuration has ATLAS site to XCache mapping.
- Rucio calculates optimal xcache node to serve the file based on site and filename.



Transfer speeds vs file size

- Follow-up to some of the questions that arose from Duncan's talk yesterday:

