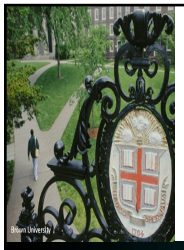

ATLAS Petascale Data Processing on the Grid: Facilitating Physics Discoveries at the LHC

Wensheng Deng [1], Alexei Klimentov [1], Pavel Nevski [1], Jonas Strandberg [2],

Junji Tojo [3], Alexandre Vaniachine [4], Rodney Walker [5]

On behalf of the ATLAS Collaboration

[1] BrookHaven National Laboratory

[2] KTH Royal Institution of Technology

[3] High Energy Accelerator Research Organization (KEK)

[4] Argonne National Laboratory

[5] LMU Munich University

Presenter: Wensheng Deng

From Collisions to Publications

Data taken by the detector go through a chain of sophisticated software

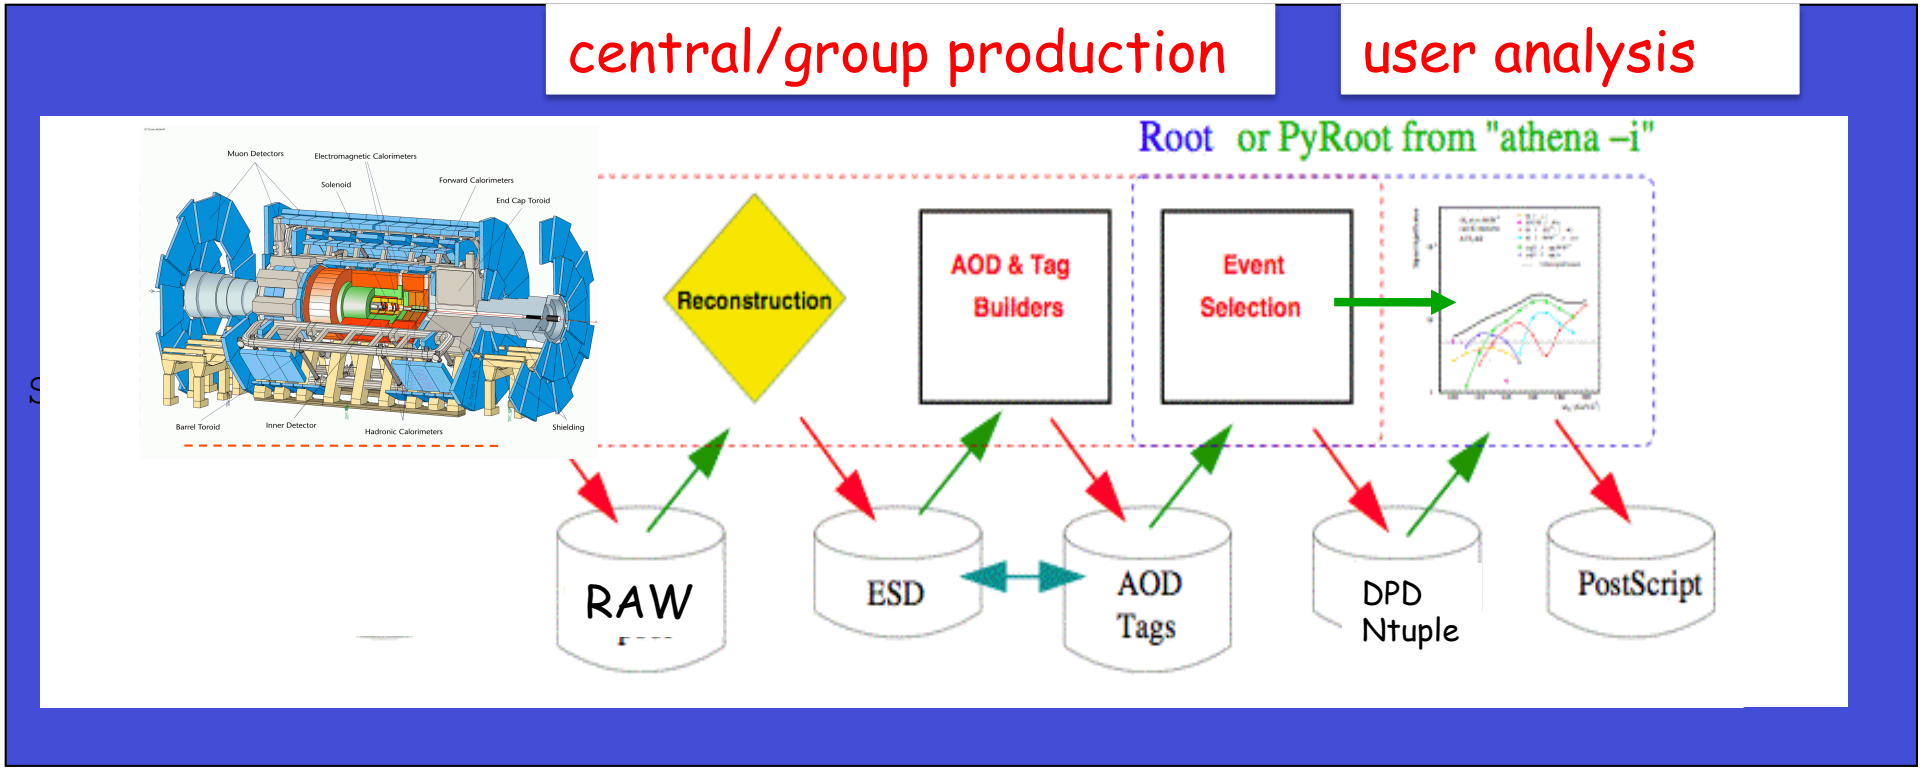

RAW Event data from TDAQ: ~ 1.6 MB
ESD (Event Summary Data): output of reconstruction (calo cells, track hits, ..): ~ 1.2 MB
AOD (Analysis Object Data): physics objects for analysis (e,γ,μ,jets, …): ~ 180 kB
DPD (Derived Physics Data): and ntuples: ~ 10-20 kB
TAG Reduced set of information for event selection: ~ 1 kB

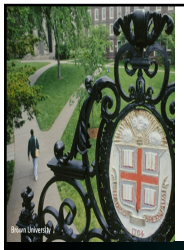

ATLAS Grid Architecture

- ATLAS runs on 3 middleware suites:
  - gLite in most of Europe and several other countries (including all A-P countries)
  - ARC in Scandinavia and a few other European countries
  - VDT in the USA
- ATLAS Grid tools interface with the middleware and shield the users from it
  - They also add a lot of functionality that is ATLAS specific
- The ATLAS Grid architecture is based on few main components:
  - Information system
  - Distributed data management (DDM)
  - Distributed production and analysis job management system (PanDA)
  - Distributed production (ProdSys) and analysis (Ganga/pAthena/prun) interfaces
  - Monitoring and Accounting tools

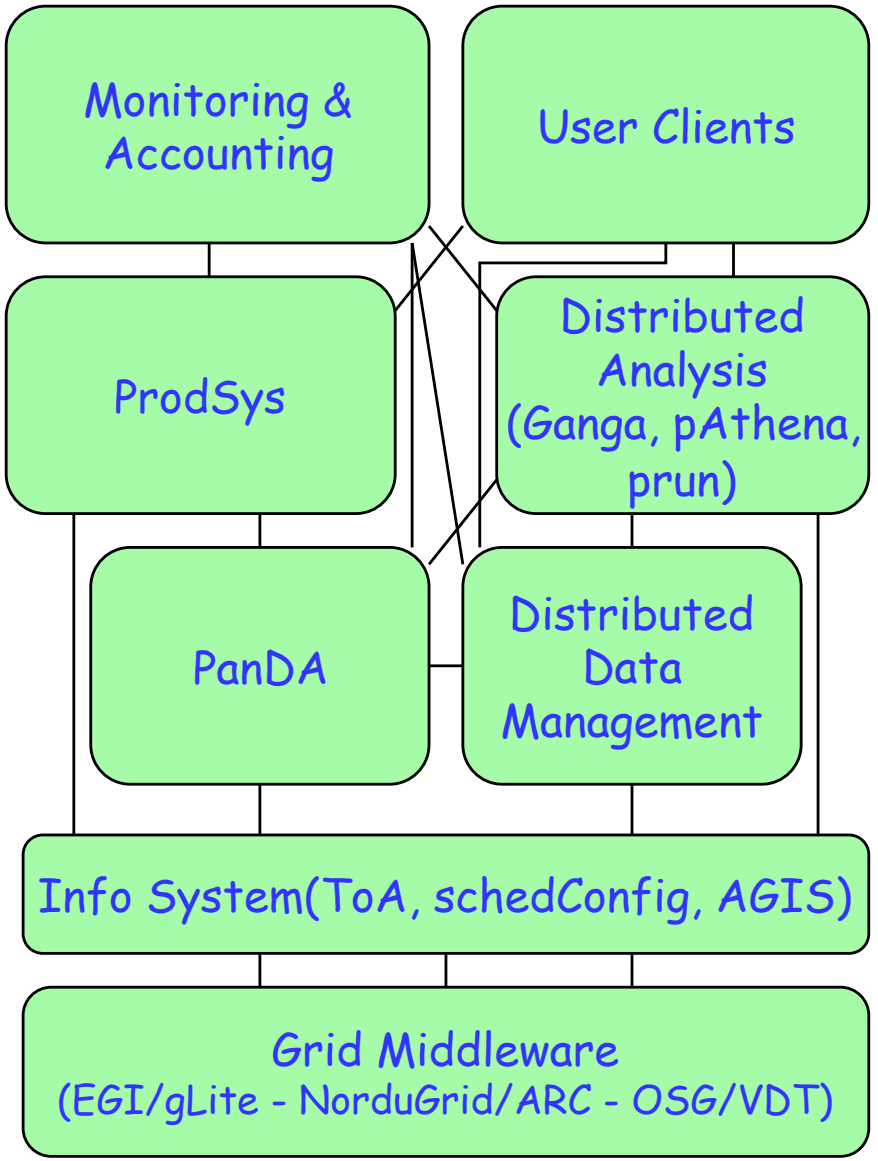

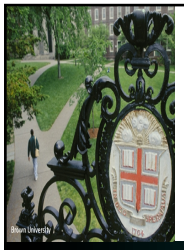

DDM/DQ2: Distributed data management

- The Distributed Data Management (DDM) architecture is implemented in the DQ2 tools and additional services

- The unit of storage and transfer is the dataset:
  - A dataset contains all files with statistically equivalent events
- DDM takes care of:
  - Distributing data produced by Tier-0 to Tier-1s and Tier-2s
  - Distributing simulated and reprocessed data produced by Tier-1/2s
  - Distributing user and group datasets as requested
  - Managing data movement generated by production activities
  - Cataloguing datasets (files, sizes, locations etc.)
  - Providing usage information for each dataset replica
  - Deleting obsolete or unnecessary replicas of datasets from disk when disks are full
  - Providing end-users with client tools to operate on datasets (import/export/move etc)

PanDA: Production and Distributed Analysis

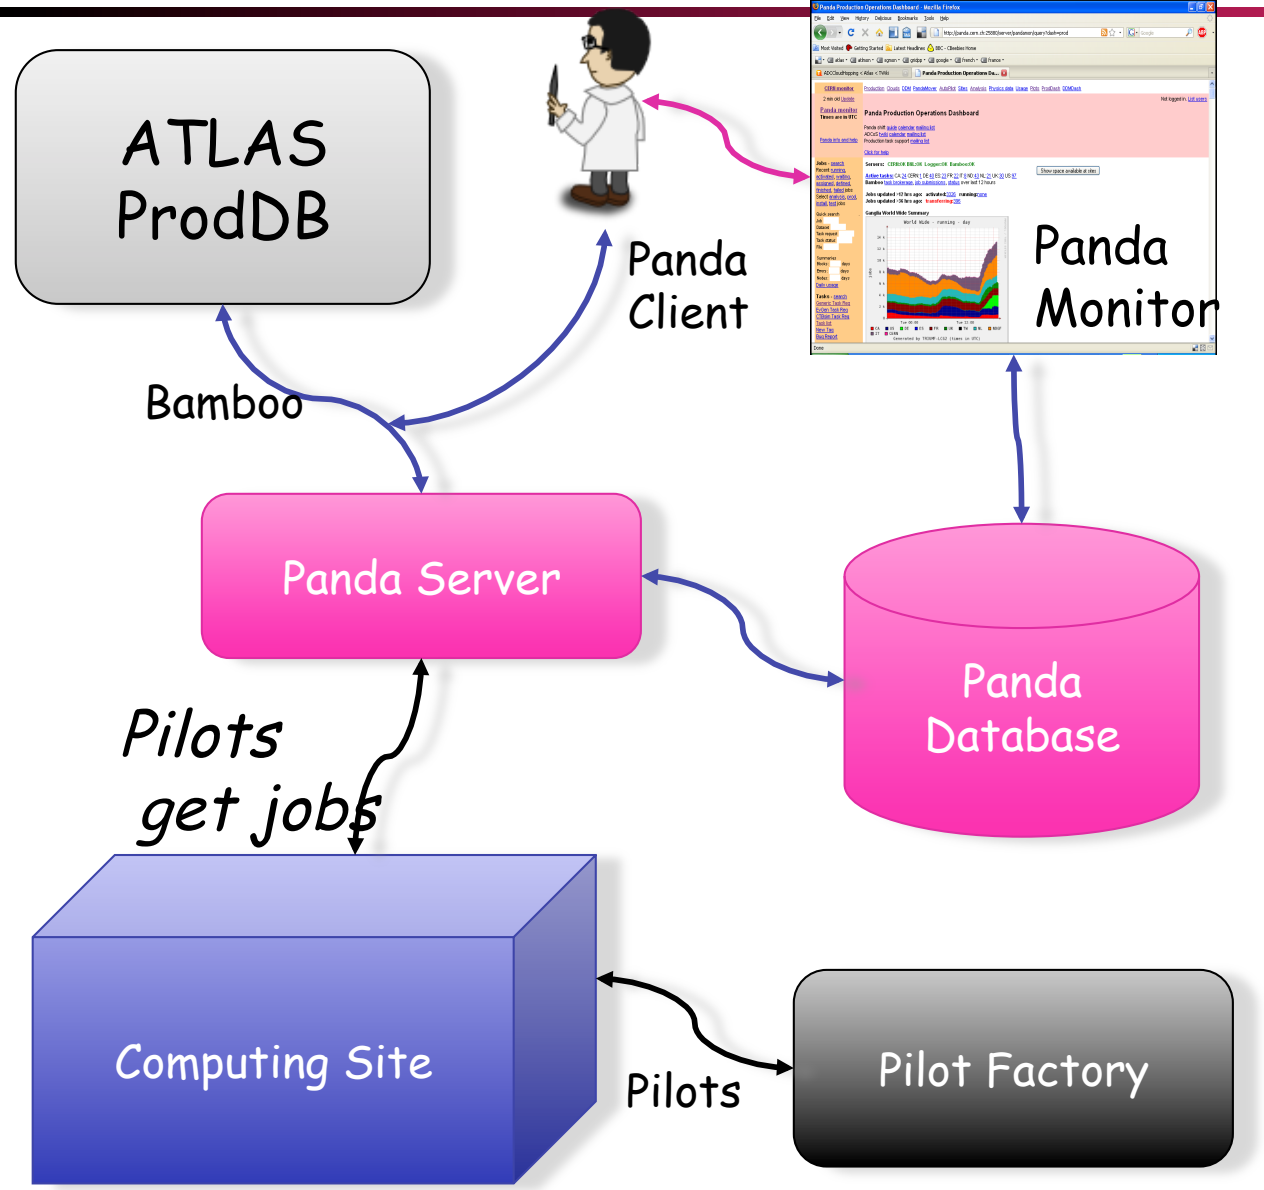

- PanDA is a pilot based workload management system
- Designed for generic applications
- Executes jobs from the production DB and from users
- Brokers jobs to sites based on available resource and data
- Ensures complex data processing on tens of thousands of CPU-cores
- Monitors site performance and supports dynamic workload sharing minimizing the processing duration

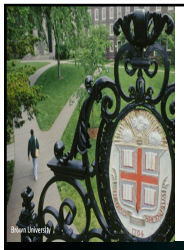

Data processing and reprocessing

- The "first-pass" processing of the raw data from the ATLAS detector is done at Tier-0, which provides promptly the data for quality assessment and analysis.
- When software and conditions/calibrations data are optimized further, collaboration decides to "reprocess" a chunk of data. This usually includes several months of data taking which is at the order of petabytes.
- The reprocessing is managed by our central production team, and utilizes ATLAS computing resources distributed world wide on the Grid. It needs dedicated efforts to ensure high quality results.
- This step requires massive data access and tens of thousands of CPU-cores, and in addition it often needs rapid software update.

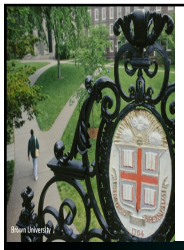

Reprocessing Goals and Requirements

Reconstruct on the grid selected large set of data with the software version signed off by Data Preparation coordination, produce and distribute bulk outputs for the collaboration analysis: DESDs, AODs, TAGs, as well as monitoring DQ histograms

- Efficient usage of computing resources on the grid

  => needs a stable and flexible production system

- Exclude possible dependence on site software or hardware

  = > assured by the Sites Validation

- Prevent bottlenecks in large-scale data access to conditions data

  => assured by the DB Release

- Exclude site-dependent failures, like unavailable resources etc

  => Strict job abortion policy

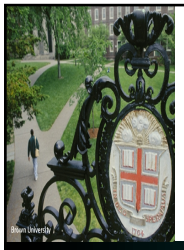

Sites Validation

- Sites Validation ensures that all sites produce numerically identical outputs for the same inputs
  - Reconstruct representative set of inputs (mixture of streams)
  - Dump all ESD, AOD collections via Athena into an ASCII files
  - Compare ASCII files to check for numerical differences
- The Sites Validation detects reconstruction and/or conditions data issues:
  - Software algorithm instabilities and/or bugs that slipped through the SW Release validation due to a limited scope of validation dataset and limited choice of sites
  - Discrepancies due to local hardware and/or system differences, such as CPU (AMD/Intel), system libraries, batch system settings, etc.
- Sites Validation is performed regularly before and during reprocessing

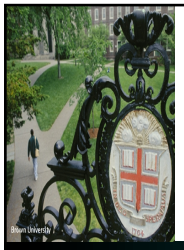

Conditions Data Access in reprocessing

- To prevent scalability problems of the database access we developed a DB Release technology, combining file-based and database data access with a software-like release technology (tarball)
- Built upon the experience with earlier reprocessing cycles, the flexible Conditions DB Release technology integrates both Conditions and Geometry data in a single dataset

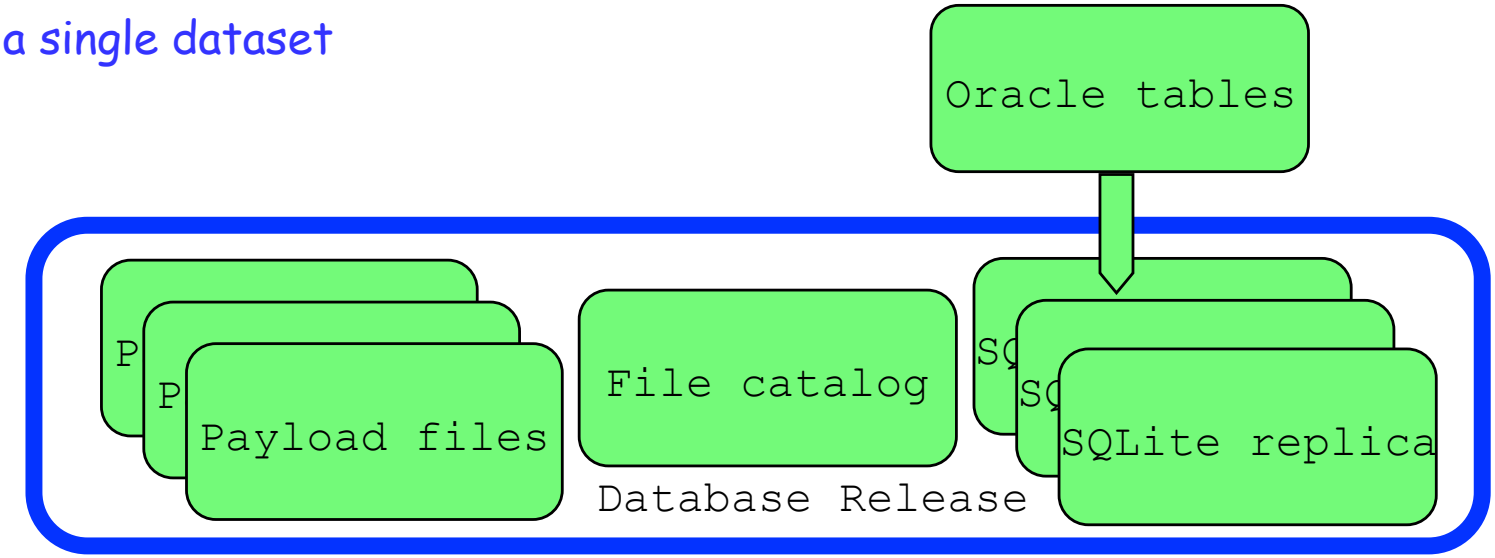

- The extensive usage of SQLite replicas for most of database-resident data minimized Oracle server load

Job Definition

date	ALL	CA	DE	ES	FR	IT	ND	NL	TW	UK	US
11-03-15	13752	1084	1455	1446	1381	2436	418	1548	38	1	3945
11-03-16	50366	2928	7470	4719	6279	4997	4592	3365	2206	926	12884
11-03-17	55238	2669	13640	3802	5621	3961	3882	3816	4418	5655	7774
11-03-18	55108	2889	12863	4159	4035	2284	4017	5403	4271	4998	10189
11-03-19	42547	732	6830	4260	2504	4007	3727	4153	4210	1825	10299
11-03-20	42160	1357	3960	5005	787	3767	3550	3032	3356	3311	14035
11-03-21	12154	913	2142	1172	11	1561	944	1034	829	1349	2199
total done	271325	12572	48360	24563	20618	23013	21130	22351	19328	18065	61325
total jobs	331953	23305	50091	25022	20632	39910	23575	34761	29526	23687	61444
% done	81.7	53.9	96.5	98.2	99.9	57.7	89.6	64.3	65.5	76.3	99.8
aborted	0	0	0	0	0	0	0	0	0	0	0
% aborted	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
jobs left	60628	10733	1731	459	14	16897	2445	12410	10198	5622	119
% left	18.3	46.1	3.5	1.8	0.1	42.3	10.4	35.7	34.5	23.7	0.2
running	11578	1178	1731	459	14	2010	2445	1284	845	1493	119
days left	-	5	0	0	0	5	0	3	3	2	1

- A key component of the production system works smoothly thanks to the efforts invested over the years

ATLAS Reprocessing Campaigns

- A year of petascale Grid Data Processing operations provides solid foundation for the ATLAS reprocessing campaign of 2011 data

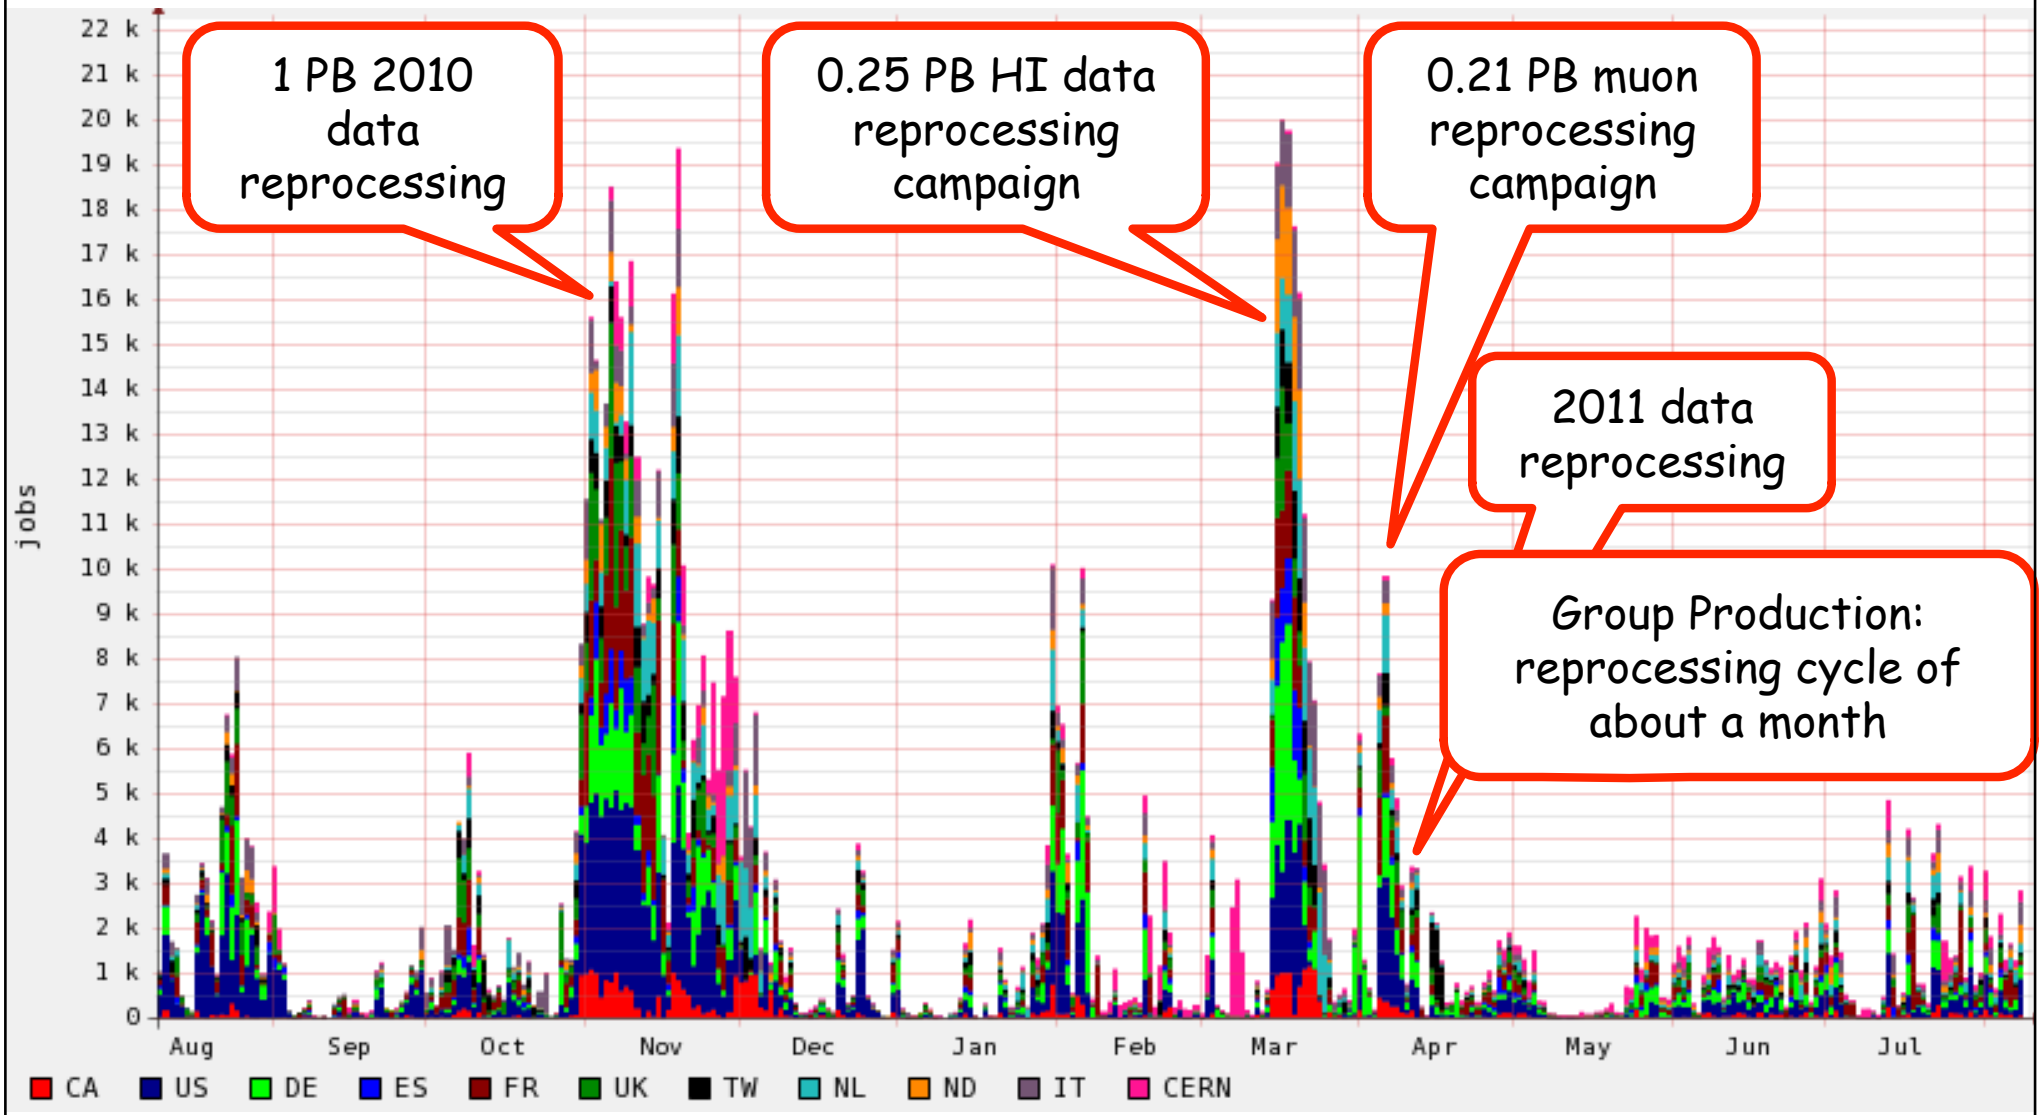

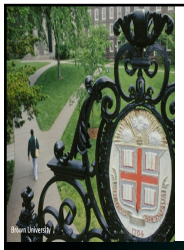

Minimize event losses

- Physics discovery requires minimal unrecoverable event losses, which are assured by an acceptable number of automatic resubmission of the failed jobs to exclude transient failures
- The events that were not possible to reconstruct during the reprocessing campaign are recovered shortly in a dedicated post-processing step using an updated software release and/or conditions
- Correcting silent data corruption in a distributed event store in peta scale is prohibitively costly. The data corruption must be detected at the spot to assure scalability.

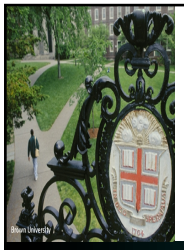

Group Production

- Reprocessing results are input to additional physics specific treatment by Physics working groups. There is a growing effort to organize data processing for physics groups in the most efficient way
- Group Production is not limited to a common version of the software:
  - Dedicated group production releases are available and centrally supported
- Our Grid Production system provided the following benefits:
  - Access to computing resources world-wide
  - Guaranteed reproducibility of results
  - Full integration with Data Management system allows automated data delivery to the final destination
  - Organized 24/7 support in case of problems
  - Flexible usage of group resources (mainly disk storage)
  - Automatic resubmission of failed jobs minimized fraction of events lost
- Common coordination was put in place between physics groups and ATLAS Distributed Computing to support group production

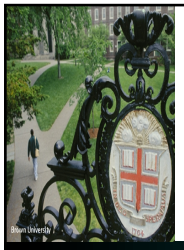

On-going activities

- Successful RAW compression validation on the Grid
  - on almost all ATLAS Tier1s.
  - large scale tests demonstrated an increased fault tolerance with compressed RAW files respect to uncompressed RAW
    - E.g. in case of transient data corruption the compressed RAW is unreadable and the job will fail but the users will not read bad data
- Adopt the Grid production system for HLT Trigger Reprocessing
  - HLT trigger community needs to validate changes before putting them online
    - e.g. new trigger release, new menu, new conditions, bug fixes
  - Using the Grid production system to go beyond the Tier0/CAF capacities
  - Test proved that the system could cope with this activity
    - HLT trigger reprocessing schedule: in full production by the end of summer

Production jobs in six months

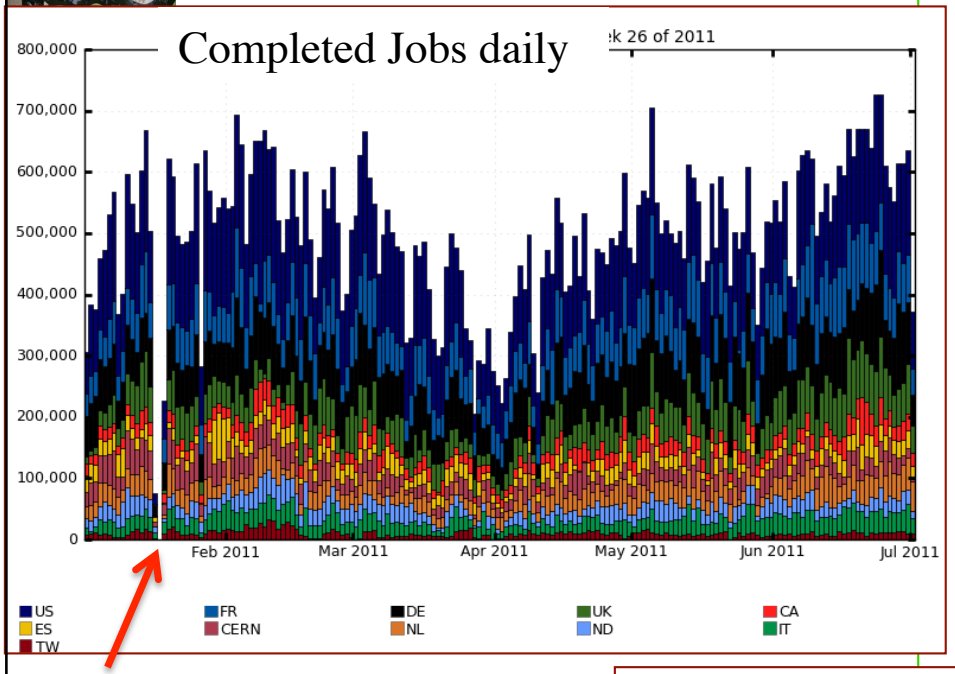

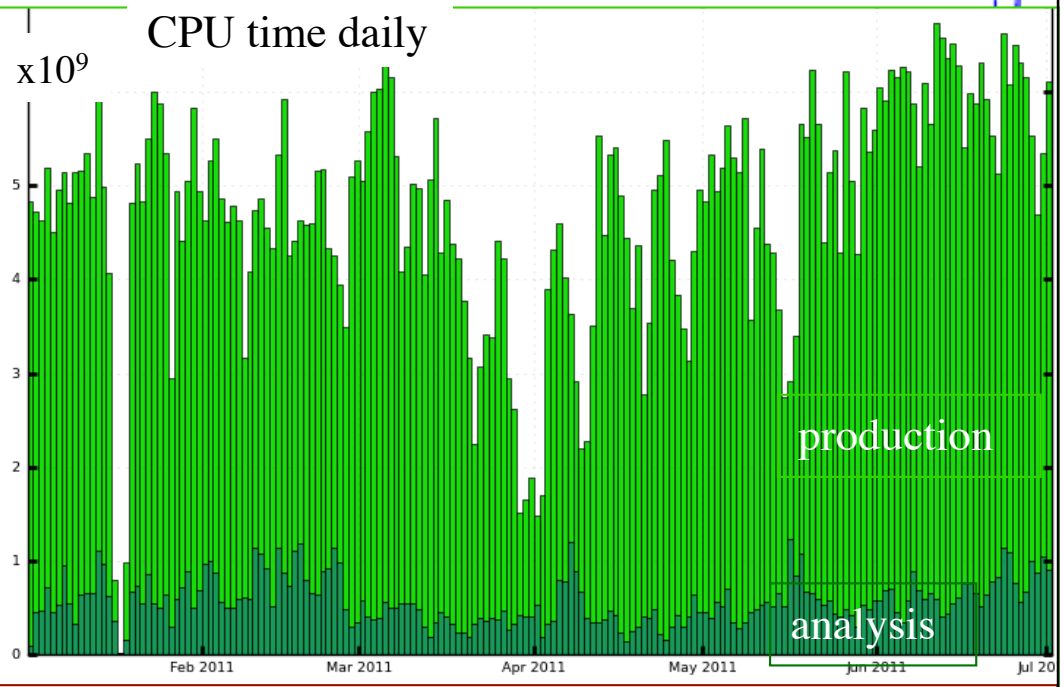

Grid global downtime to scale up DB capacity for >=2011

- CPU ~ 4 10^9 sec daily
- > 600K jobs /day
  - ~ 80k jobs in parallel

- Reprocessing of 0.5 PB of select data: HI, Muon, Fast, Tags
- ~ one cycle/month of group reprocessing (NTUPs)

MC production 2011

- new G4 production: 8400 tasks: 800M events produced
- digi + reco
  - MC10a campaign (2011 data analysis)
  - MC10b campaign (2011 data analysis)
- MC11 new event generation and G4 started: progressing as expected

Reprocessing of 2011 Data

- Following thorough preparations, ATLAS is ready to start with the bulk reprocessing for the most of 2011 data (about 1 PB)

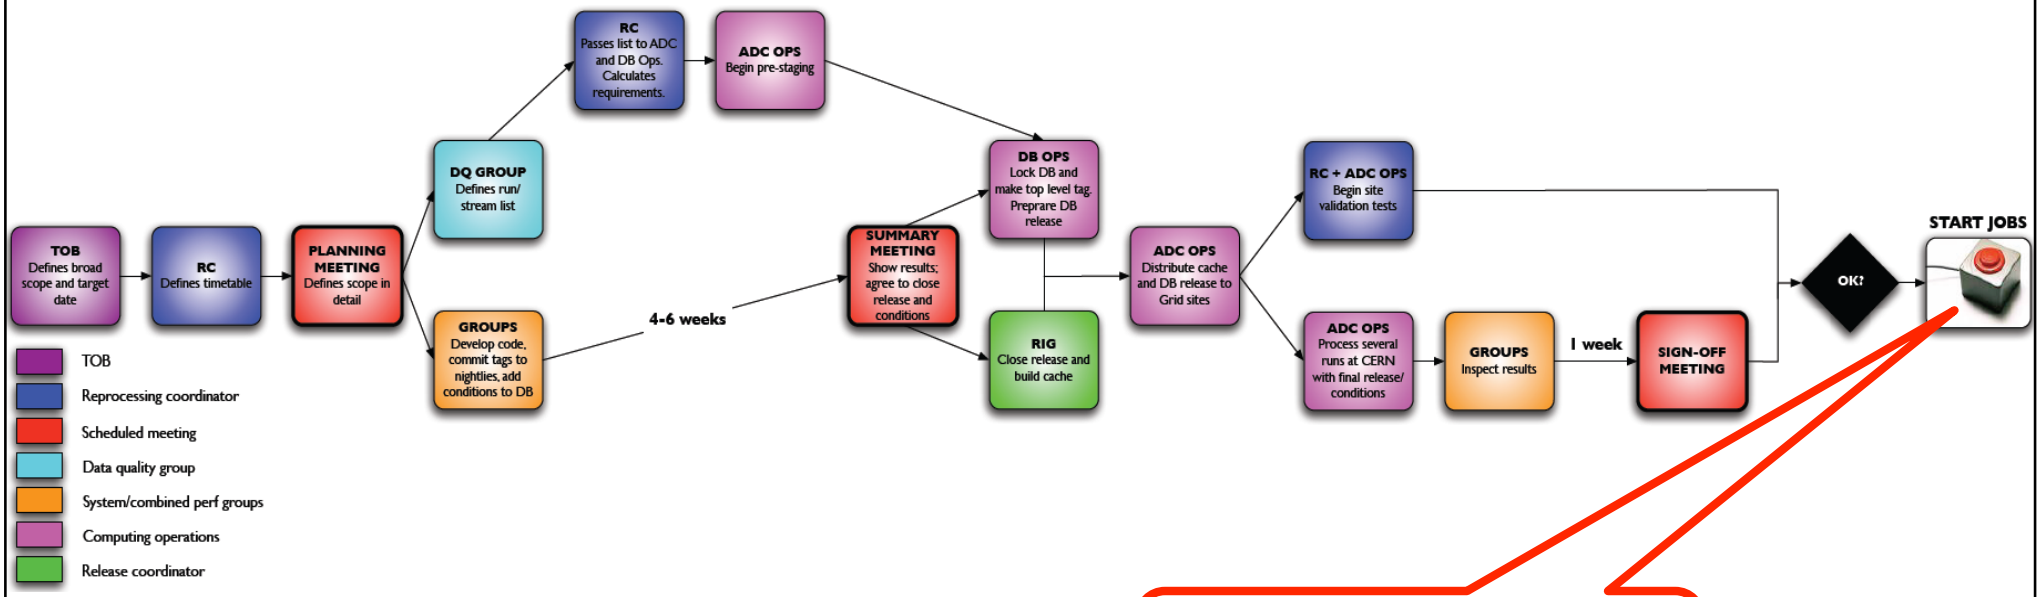

Started today!

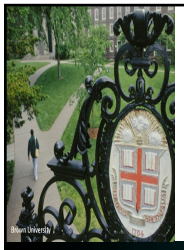

Conclusions

- Our Grid production system coped well with LHC data
- Since the start of LHC data taking ATLAS has conducted many successful reprocessing campaigns on the Grid
- In addition production system empowered many physics groups with coordinated access to the Grid. Group Production consumed most of Grid CPU resources
- ATLAS grid production is capable to reach high efficiency and reliability ready for the petascale challenges of the LHC long run