ATLAS distributed computing: experience and evolution

A Nairz[1], on behalf of the ATLAS Collaboration

[1] CERN, CH–1211 Geneva 23, Switzerland

E-mail: armin.nairzcern.ch

Abstract. The ATLAS experiment has just concluded its first running period which commenced in 2010. After two years of remarkable performance from the LHC and ATLAS, the experiment has accumulated more than 25 fb^{-1} of data. The total volume of beam and simulated data products exceeds 100 PB distributed across more than 150 computing centres around the world, managed by the experiment's distributed data management system. These sites have provided up to 150,000 computing cores to ATLAS's global production and analysis processing system, enabling a rich physics programme including the discovery of the Higgs-like boson in 2012. The wealth of accumulated experience in global data-intensive computing at this massive scale, and the considerably more challenging requirements of LHC computing from 2015 when the LHC resumes operation, are driving a comprehensive design and development cycle to prepare a revised computing model together with data processing and management systems able to meet the demands of higher trigger rates, energies and event complexities. An essential requirement will be the efficient utilisation of current and future processor technologies as well as a broad range of computing platforms, including supercomputing and cloud resources. We will report on experience gained thus far and our progress in preparing ATLAS computing for the future.

1. Introduction

The first extended operations period of the Large Hadron Collider (LHC) from 2010 to early 2013 (the so-called "Run-1") was extremely successful for the ATLAS experiment [1]. In total, ATLAS recorded an integrated luminosity of about 27 fb^{-1} of pp collision data (5.3 fb^{-1} at a centre-of-mass energy of $\sqrt{s} = 7$ TeV in 2010–2011, 21.7 fb^{-1} at $\sqrt{s} = 8$ TeV in 2012), corresponding to 7.6 billion events and a data volume of 7.4 PB, as illustrated in figures 1 and 2. The greatest physical success was undoubtedly the discovery of a Higgs-like boson, as announced in July 2012.

For about one month per annual run period, the LHC was colliding heavy ions (Pb-Pb in 2010–2011, p-Pb in 2013); during those periods ATLAS recorded 167 pb^{-1} of Pb-Pb and 30 nb^{-1} of p-Pb collision data, amounting in total to 720 million events and a data volume of 740 TB. Based on this data sample, ATLAS has been pursuing a rich, fruitful research programme on exploring the properties of the quark-gluon plasma.

The ATLAS distributed computing system, the subject of this paper, was also working extremely well during Run-1 [2],[3]. It has proven an indispensable component of the ATLAS experiment and contributed substantially to its overall success.

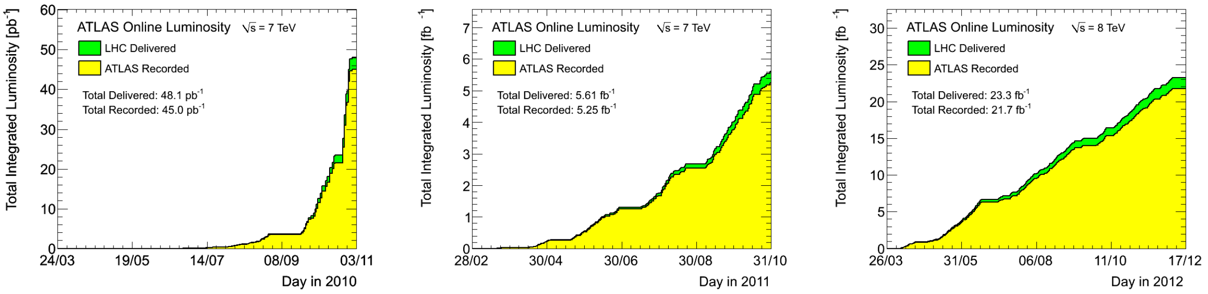

Figure 1. Integrated luminosity of *pp* collisions data delivered by the LHC (green histograms) and recorded by ATLAS (yellow histograms) in 2010 (left), 2011 (centre), and 2012 (right).

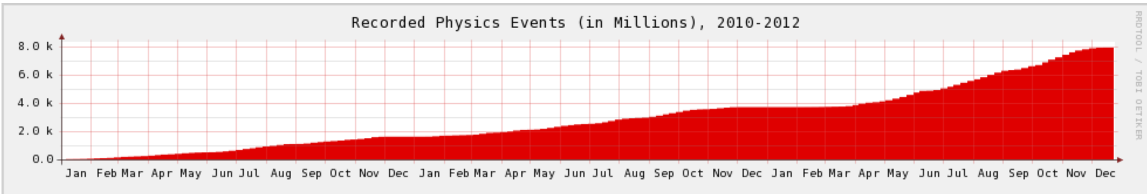

Figure 2. Number of *pp* collision events for physics analysis ("physics events") recorded by ATLAS in the period 2010–2012.

This paper is organised in two parts. The first part introduces the ATLAS computing system (section 2) and describes its main constituents, Tier-0 (section 3), distributed data processing and data management (section 4), together with their respective performances and achievements during Run-1. The second part (section 5) discusses the challenges ATLAS computing has to face for the next LHC operation period ("Run-2"), and gives a short overview of the strategies adopted and projects launched to address them, in order to get prepared for the restart of data taking in 2015.

2. Overview of the ATLAS computing system

The ATLAS Distributed Computing (ADC) system manages the worldwide data processing, Monte Carlo simulation production and user analysis activities, together with their associated data accesses and transfers. There are over 100 computing centres (sites) participating in the ADC project; the available, managed resources (as of May 2013) amount to up to 150k computing cores, about 70 PB of total disk space, and about 50 PB of tape-resident data.

2.1. The ATLAS computing model

The sites participating in ADC are part of the more comprehensive LHC Computing Grid (LCG) [4] and correspondingly organised in a hierarchical ("tiered") structure, according to their roles, available infrastructure and resources:

- the Tier-0 centre at CERN, which is mainly responsible for recording raw data onto tape, carrying out calibration processing and first-pass physics data processing;
- ten Tier-1 centres (i.e. large computing centres distributed across the world and well-connected with the Tier-0), whose tasks are the storage of a replica of the raw data on tape and data reprocessing;
- several dozens of Tier-2 centres (i.e. regional computing centres providing disk and CPU resources and associated to their respective Tier-1 centres), which are responsible for running group and end-user analysis jobs.

Monte Carlo simulation production is run both on Tier-1s and Tier-2s, wherever possible. Data are distributed and replicated over the Grid, both for redundancy (i.e. to secure data with replicas) and for increasing the accessibility (i.e. to create more replicas for frequently used, popular data).

Starting from raw data, as recorded by the ATLAS detector, processing at the Tier-0 and on the Grid produces a set of derived data products, serving dedicated purposes:

- ESD (Event Summary Data): output objects of reconstruction
- AOD (Analysis Object Data): an event representation (ATLAS-wide format) with reduced information for physics analysis
- DPD (Derived Physics Data): a data representation for end-user analysis. DPDs are produced for working groups and individual end-users and could have a group-specific format. Examples of DPD formats are dESD (derived from ESD, for performance groups), dAOD (derived from AOD, for physics groups) or NTUP (ROOT n-tuples, both for physics groups and end-users).
- TAG: a format holding event-level metadata (event tags; i.e. short event summaries primarily used for event selection)

A schematic view of the raw and derived data types, together with the associated workflows to produce them, is shown in figure 3.

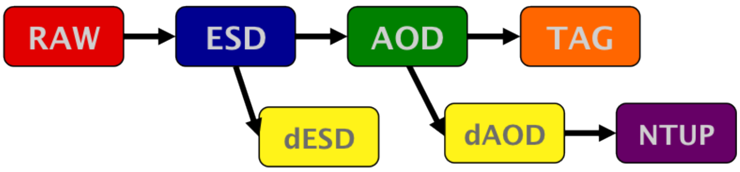

Figure 3. Schematic view of ATLAS raw and derived data types and associated workflows.

More details about the ATLAS computing model, as originally laid out in 2005–2008, can be found in references [5],[6].

2.2. Evolution of the ATLAS computing model

During Run-1 it turned out necessary to adjust the original computing model (as sketched in section 2.1) to new use cases, changing run conditions and technology advances [7],[8].

In response to unexpected usage patterns for the individual data types, data distribution policies were revised and adjusted; e.g. it turned out that users preferred the variety of DPD formats over AODs for their analyses, thus the number of AOD replicas over the Grid could be reduced. In addition, a more dynamic data placement strategy according to data popularity was introduced.

Also run and data-taking conditions changed in the course of Run-1: upgrades in the ATLAS data-acquisition infrastructure allowed an increase of the event recording rate (from originally foreseen 200 Hz to eventual 400–500 Hz), and advances in the LHC led to higher luminosity, resulting in higher pile-up and bigger event sizes. Through introduction of offline raw-data compression at Tier-0, the raw-data volume could be significantly reduced (almost by a factor two), thus saving tape space and allowing for the longer-term placement of raw data on Tier-1 disks to facilitate the access. Since ESD event sizes and their total data volume on Tier-1 disk grew too big, it was decided to keep bulk ESDs on disk only for a limited lifetime (typically a few months) and delete them afterwards.

As the network connectivity between sites had been improving significantly over the years, the original static and hierarchical organisation of Grid sites [9] and the data transfer policy between them was revisited. Based on the measured network performance, ADC moved gradually from a “tree-like” topology with Tier-1/Tier-2 association to a “mesh” of any Tier-$*$ combination, allowing direct data transfers between all kinds of sites. This resulted in faster transfers, a more efficient use of the computing infrastructure with less load on the system, and eventually more end-user convenience.

3. Tier-0 processing

The ATLAS Tier-0 system (see [10] for a more detailed system overview) was running very reliably, stably and successfully during the first LHC data-taking period.

First-pass data processing at Tier-0 was keeping up with the increasing LHC performance. Dedicated computing resources, both storage and CPU, got adjusted and extended every year to meet the requirements. In 2012 Tier-0 resources comprised more than 2 PB of disk (with tape back-end) and about 3k computing cores of the CERN batch system. The CPU capacity could be increased on demand by up to 4.5k cores through flexible usage of additional, shared public batch resources ("spill-over" mechanism). Figure 4 shows, for the 2012 data-taking period, monitoring plots of the physics raw data rate, as recorded by the ATLAS data-acquisition system and transferred to the Tier-0 (with peaks of 1.3 GB/s), and of the number of concurrently running Tier-0 jobs. (Note that for most of the period Tier-0 was running about 6k parallel jobs. A peak activity of about 7.5k parallel jobs is clearly visible in June 2012, an intensive processing period preceding the discovery of the Higgs-like boson.)

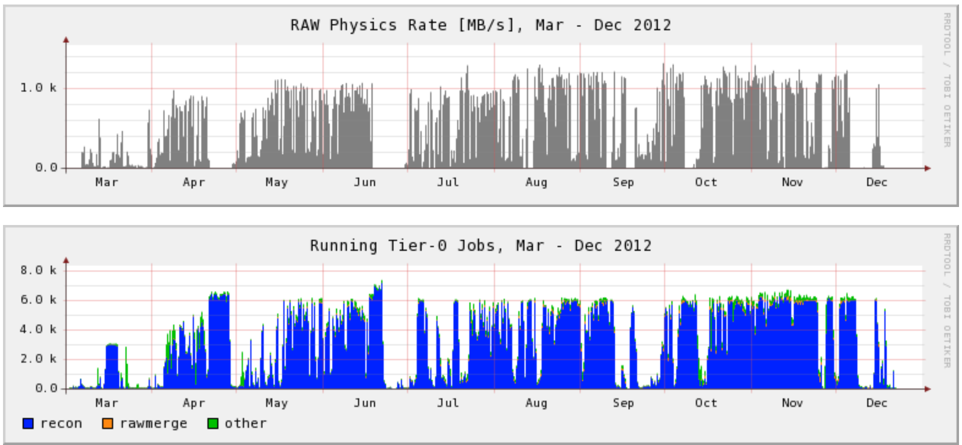

Figure 4. Recorded ATLAS physics raw data rate (top) and concurrently running Tier-0 reconstruction jobs (bottom) for the 2012 data-taking period.

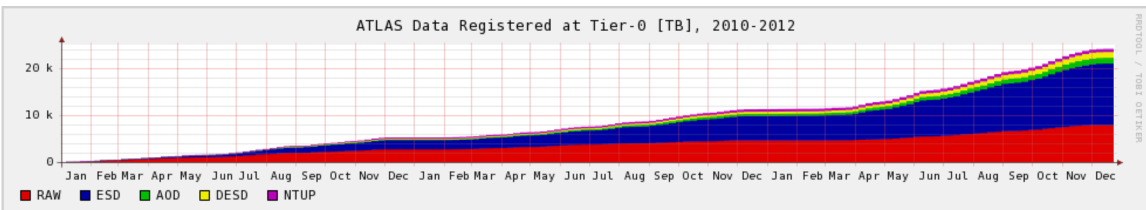

Figure 5. Accumulated total volume of ATLAS raw and derived data types processed by Tier-0 and registered with the DDM system in the period 2010–2012.

Tier-0 processing was organised in two separate steps: in a first step, the so-called "calibration loop", dedicated smaller data streams (most notably the "express stream") were processed with high priority, and the resulting output was provided to detector and data-quality monitoring groups for deriving calibration and alignment information and assessing the data quality. After a delay of usually 48 hours, the second step – processing of the bulk physics data – was launched, already using up-to-date calibration. This resulted in high-quality data reconstruction already from first-pass processing, such that most 2012 data produced by Tier-0 could be used directly in physics analysis, without need of prior reprocessing.

Figure 5 shows the accumulated data volume processed by Tier-0 and registered with the ATLAS distributed data management (DDM) system, thus prepared for the export to the Grid and made available to the users.

4. Distributed computing

The ATLAS distributed computing (ADC) system and infrastructure were working very well and successfully during Run-1 [2],[3].

The ATLAS production system is responsible for managing data processing and running user analysis jobs on the Grid; it is based on PanDA (see [11] for more details) as workload management system. During peak activity in 2012, the system was running up to 150k parallel jobs (figure 6); the usual average job throughput (i.e. the number of successfully completed jobs) per day amounts to about one million jobs, with peaks of up to 1.7 million (figure 7).

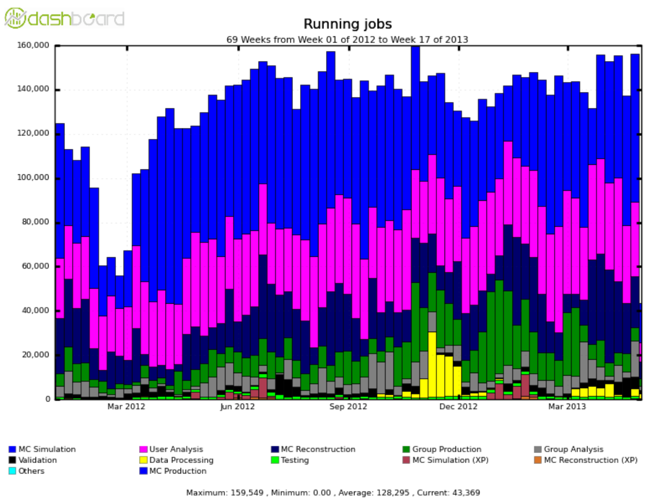

Figure 6. Number of concurrently running ATLAS Grid jobs per activity for the period from January 2012 to April 2013.

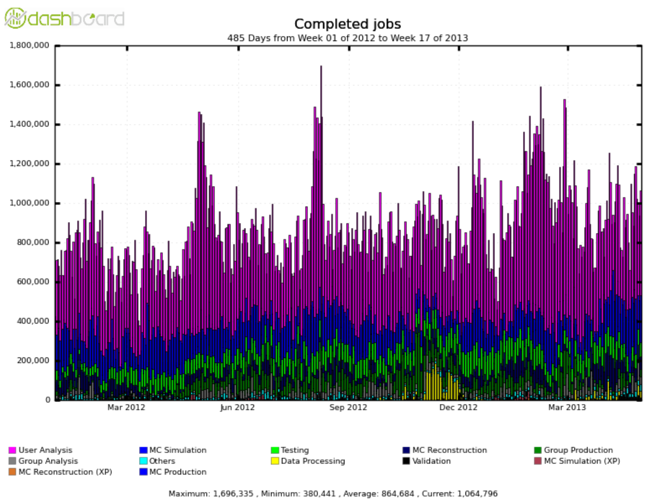

Figure 7. Number of successfully completed ATLAS Grid jobs per day and activity for the period from January 2012 to April 2013.

The majority of jobs run on the Grid are short end-user analysis jobs. Submission of analysis jobs to PanDA takes place through a powerful and simple command-line interface. User-friendly and equally easy-to-use data management tools allow users to have the job output either transferred to their "home" site on the Grid (both automatically and manually), or downloaded to off-Grid computers.

The ATLAS distributed data management (DDM) system (see [12],[13] for more details) manages the data distribution to and between Grid sites. The system successfully completes about one million file transfers per day, amounting to a total daily volume of 400–800 TB (figure 8); those numbers include transfer of the input/output of both production and user-analysis jobs.

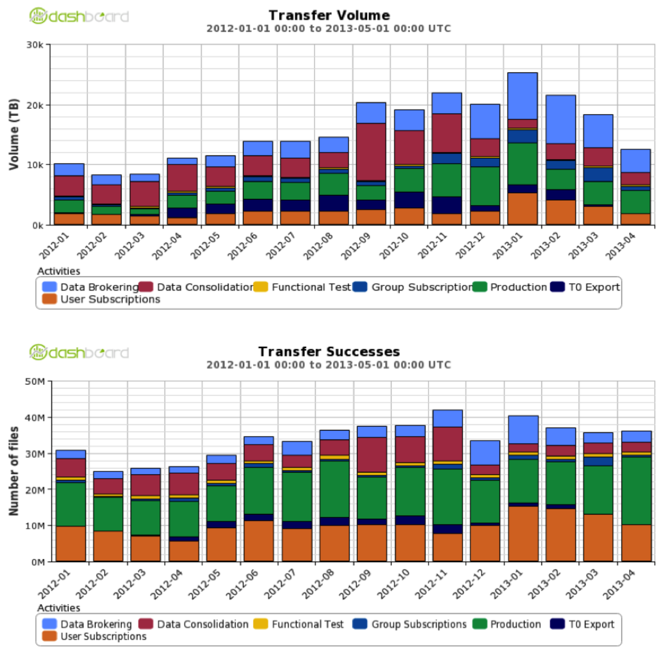

Figure 8. Monthly data volume (top) and number of files (bottom) transferred by the ATLAS DDM system over the Grid for the period from January 2012 to April 2013, for different activities.

As illustrated through figures 6–8 (examples only!), ADC is endowed with a comprehensive monitoring suite [14], which is essential for day-to-day operations, provides the tools for shift services, or allows to extract overview and resource usage summaries.

5. Future challenges: Run-2 and beyond

Regardless of the success and achievements of ATLAS distributed computing during Run-1, conditions for Run-2 (starting in 2015) will be much more demanding, both in terms of external run conditions (LHC, ATLAS data acquisition) and available computing resources.

Higher LHC centre-of-mass energy (13–14 TeV) and higher luminosity will result in higher event multiplicities and more pile-up (i.e. number of *pp* interactions per bunch crossing, or overlaid events). In 2012 the mean pile-up, $\langle\mu\rangle$, was at about $\langle\mu\rangle = 20.7$, requiring about 20–30 seconds wall-time at the Tier-0 to fully reconstruct an event. In Run-2 one has to expect values of $\langle\mu\rangle = 40$ in the most "favourable" circumstances (luminosity levelling at the LHC), but the pile-up could easily be higher. Using the same reconstruction software as in 2012, the projected reconstruction time per event, for a centre-of-mass energy of 14 TeV and $\langle\mu\rangle = 40$, would be at about 70 seconds (figure 9), i.e. more than twice as much as in 2012.

There are also hardware and software upgrades planned in the ATLAS online trigger and data-acquisition system, in order to achieve higher data-taking rates; the target figure is 1 kHz (sustained), i.e. about twice the rate as in 2012.

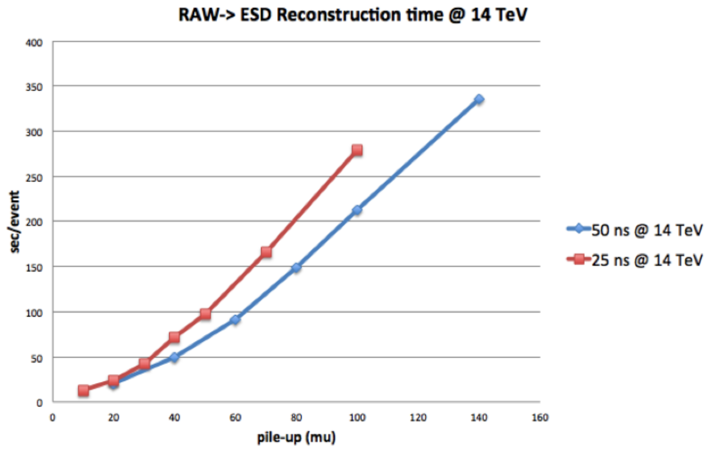

Figure 9. Projected reconstruction time for a *pp* collision event at a centre-of-mass energy of 14 TeV, dependent on the pile-up, for two LHC operations scenarios: a bunch-crossing rate of 40 MHz (25 ns bunch spacing; red graph) and of 20 MHz (50 ns bunch spacing; blue graph), respectively.

To account for the changes in LHC run and ATLAS data-taking conditions, more computing resources will be needed everywhere: CPU, storage (due to more data and bigger event sizes), network, etc. Realistically, however, due to a stretched financial situation, ATLAS cannot expect an increase in the computing budget; facing a "flat budget" scenario, an increase in Grid resources can only be achieved through technological progress. Therefore, in general, the goal will be to better utilise available resources and explore new "affordable" resources wherever possible.

To address those issues and increase the overall efficiency of the system, the ATLAS software and computing community has launched an ambitious upgrade and development programme, which will be introduced briefly and non-exhaustively in the following sections.

5.1. New DDM system (Rucio)

As mentioned in section 4, the current ATLAS DDM system (DQ2) [12],[13] has been working very well and lived up to (and even beyond) expectations. Currently (May 2013) the system administers about 140 PB of ATLAS data on more than 150 Grid sites, and the total data volume typically increases per year by more than 40 PB in more than 100 million files; see figure 10 as an illustration.

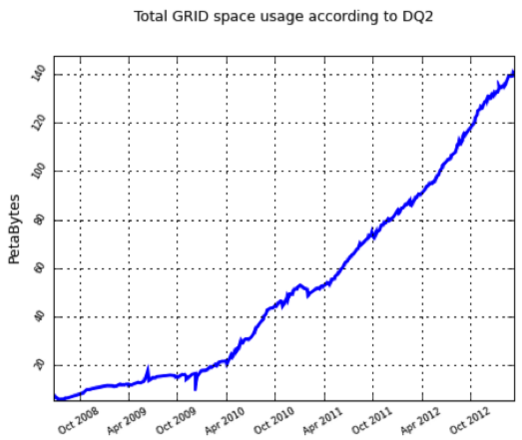

Figure 10. Total volume of data managed by DDM in the period from July 2008 to March 2013. Note the sharp increase since April 2010, the beginning of the 2010 data-taking period.

Nevertheless, projections towards Run-2 conditions indicate that DQ2 will not continue to scale. There are matter-of-principle extensibility issues, due to its design, together with changes to the computing model and the Grid middleware, that are difficult to accommodate.

The ATLAS DDM developers' team has therefore embarked on the project of a fundamental redesign and rewrite of the DDM software (project name *Rucio*) [15]. The next-generation DDM system will address the above-mentioned scalability issues. It will have new concepts for optimisation of storage occupancy, data distribution and network utilisation, and will provide better means of accounting and space quota management. Where appropriate, Rucio will be exploiting new technologies; e.g. it envisages the usage of noSQL database technologies (like Hadoop) and foresees the support for new protocols beyond *gridftp* (like *xroot* and *http*). Rucio development is planned to be concluded by the end of 2013; a fully functional new system should then be ready to go in production in early 2014.

5.2. New Grid production system (ProdSys2)

As mentioned in section 4, the existing Grid data processing system [11] has been working very well. There are, however, similar scalability concerns as with DQ2 (section 5.1), together with difficulties to implement new features and requirements in the existing design.

Again, a fundamental redesign and rewrite campaign for a new Grid production system (project name *ProdSys2*) is in progress [16]. The new system will address the above-mentioned scalability concerns. It will provide better support for new or different workflows, and will offer means of more flexible and dynamic job definitions and job scheduling, e.g. support for multi-core node scheduling and dynamic job splitting. ProdSys2 is planned to be tightly integrated with the "Event Service", a newly conceived service being developed right now, which will enable data processing at a much finer granularity (at event level, instead of the currently only possible file level). Wherever beneficial for the efficacy of the ProdSys2 system, new Rucio concepts (e.g. file-based data management, usage of transfer rules) are planned to be exploited. Similar to Rucio, ProdSys2 is planned to be deployed and to replace the existing system in early 2014.

5.3. Concurrency

Much effort, in the whole community, is presently invested into making better and more efficient use of modern multi-core processors. The issue has also been in the focus of the ACAT2013 workshop (see e.g. [17]).

The ATLAS software and computing community is following several tracks to address the issue. Firstly, effort is going into optimising the existing code [18], parallelisation of algorithms and evaluation of new, customised software libraries. Secondly, work is ongoing on a new concurrent framework, which is planned to support features like full threading and event-level/algorithm-level parallelism. In this area, ATLAS is greatly benefitting from the collaboration with other CERN teams. And thirdly, ATLAS is pursuing the development, testing and deployment of *AthenaMP* [19] (*Athena*: the ATLAS core software package, *MP*: MultiProcess), a framework featuring basic parallelism through processing the event loop by spawning parallel worker processes. AthenaMP also has the advantage of reducing the memory footprint of a job through memory sharing ("copy-on-write" feature) between workers.

AthenaMP is already advanced and will be commissioned and moved to production in the course of 2014. Some operational challenges on the Grid still have to be addressed and resolved [20], such as the efficient matching of jobs to resources (which involves issues like whole-node scheduling, the deployment of multi-core queues, or the optimal usage of sites offering serial and multi-core queues). Flexible input/output file management in AthenaMP is also mandatory and needs to be refined, in particular once the new Event Service (section 5.2) has become operational.

5.4. Cloud computing

In the following two sections ADC initiatives will be introduced whose common goal is the extension of the CPU capacity through the usage of "opportunistic resources". One possibility is the usage of free (e.g. academic) clouds, or alternatively, "cheap" commercial clouds as temporary ("opportunistic") extensions of the ATLAS CPU resources [21]. Cloud computing also has been intensively discussed during the ACAT2013 workshop (see e.g. [22]).

The utilisation of statically allocated virtual machines (VMs) in a cloud has already been successfully demonstrated in the context of ATLAS distributed data processing. Examples are the adaptation of the ATLAS high-level trigger farm for running simulation production during the period of the long LHC shutdown, or several cloud-based PanDA queues that have been running production jobs. In that context, CERN OpenStack as framework and top-level infrastructure has been commissioned and is being used, offering additional 800 computing cores to ATLAS since October 2012.

In addition, ADC is already actively exploring public-private partnerships in cloud computing. This is still an R&D project, but ATLAS has managed to get access to such resources on a scale of several thousand cores, and to run its jobs over sustained periods of weeks to months. As a pioneering activity, ATLAS actively participated in the Helix-Nebula project; recently new projects were started in collaboration with commercial cloud providers, like Amazon (EC2) or Google (GCE).

There are still many open issues to address and resolve, such as the efficient usage of cloud storage, or the way transfers to/from clouds (e.g. for user analysis jobs) are handled. Overall, however, cloud computing is a quickly evolving field offering interesting and attractive opportunities.

5.5. Usage of HPCs

Another possibility of utilising opportunistic resources (see section 5.4) is gaining access to cycles on high-performance computing (HPC) facilities. HPCs are frequently operated by academic or research institutions (or funded by the same agencies that also fund HEP research), which often facilitates the access for members of the HEP community to those facilities. HPCs are particularly attractive because of their huge capacity: equivalents of several hundred thousands of computing cores in a single HPC could easily surpass the capacity of the LCG as a whole, so gaining access to even a small fraction of an HPC may have a substantial impact.

ADC has managed to run first proof-of-principle tests on several HPCs in the US and Europe, e.g. Intrepid (Argonne, USA), SuperMUC (Munich, Germany) and MOGON (Mainz, Germany).

Several open issues regarding the efficient usage of HPCs still have to be resolved. Examples are whole-node scheduling on HPCs, or the question of how to run jobs in an environment where there is no worker-node disk and no outbound I/O. In the context of job scheduling, one has to be able to get fast onto those facilities once they become available; to get fast off them once they disappear; jobs and the job management system have to be robust against disappearing resources; and one has to be able to fill the HPCs with fine-grained workflows. All those issues will be addressed by the new ProdSys2 (section 5.2).

5.6. Federated storage

Using "storage federations" (i.e. clusters of well-connected sites sharing their storage and appearing as an entity to the outside) constitute a means of optimising storage resources, for example through reduced data transfers due to data placement. The concept heavily relies on the good connectivity between sites, and has been made possible by the fast evolution of the network infrastructure in recent years; thus, in a way, the network may become a managed resource as well. Storage federations have also received much attention during the ACAT2013 workshop (see e.g. [22]).

The traditional ATLAS computing model is based on the "DataGrid" concept, where jobs go to the data (and jobs access the data through the local network, LAN) and data (usually whole data sets) may be replicated across sites for better accessibility. Storage federations, on the other hand, provide new access modes and redundancy: jobs now access the data on shared storage resources through the wide-area network (WAN); as (e.g. analysis) jobs may not need all the input files of a whole data set, it may be sufficient to transfer only a part of the data set across sites, or the possibilities of file or event-level caching could be exploited; and, by creating a common name-space across all storage sites, an easy-to-use and homogeneous way of accessing data could be established.

ADC has started evaluating federated storage based on a system of XRootD "redirectors", which turned out to be a working solution [23]. Feasibility studies have been positively concluded after a "full dress rehearsal" exercise carried out in spring 2013, involving sites in the US and Europe. The plan is to include more sites (e.g. also in Asia) in future tests.

6. Conclusions

ATLAS distributed computing has been working extremely successfully during Run-1 in all aspects of large-scale data processing, data management and user analysis.

To get ATLAS distributed computing prepared for the challenging start-up of the LHC in 2015, many interesting and promising ideas and projects are being pursued. The overall goal is to use the scarce computing resources as efficiently as possible, and to open up new resources in an innovative way (like clouds or HPCs) wherever feasible. The projects make good progress, but there is still a lot of work ahead in the years 2013–2014.

Acknowledgements

The author wishes to thank in particular the following ATLAS colleagues for providing valuable information and material: I. Ueda, S. Campana, K. De, R. Gardner, V. Garonne, B. Kersevan, A. Klimentov, T. Lecompte, T. Maeno, S. Panitkin, M. Potekhin, H. Sakamoto, R. Seuster, I. Vukotic, R. Walker and T. Wenaus.

References

[1] The ATLAS Collaboration 2008 The ATLAS Experiment at the CERN Large Hadron Collider *J. Inst.* **3** S08003

[2] Ueda I (for the ATLAS Collaboration) 2012 ATLAS Operations: Experience and Evolution in the Data Taking Era *J. Phys.: Conf. Ser.* **331** 072034

[3] Ueda I (for the ATLAS Collaboration) 2012 The Present and Future Challenges of Distributed Computing in the ATLAS Experiment *Proc. of Science* PoS(ICHEP2012)500

[4] Eck C et al. 2005 *LHC computing Grid: Technical Design Report* CERN-LHCC-2005-024 (Geneva: CERN)

[5] The ATLAS Collaboration 2005 *ATLAS Computing: Technical Design Report* CERN-LHCC-2005-022 (Geneva: CERN)

[6] Jones R and Barberis D 2008 The ATLAS computing model *J. Phys.: Conf. Ser.* **119** 072020

[7] Jones R W L and Barberis D 2010 The evolution of the ATLAS computing model *J. Phys.: Conf. Ser.* **219** 072037

[8] Campana S, Barreiro Megino F, Jézéquel S, Negri G, Serfon C and Ueda I 2012 Evolving ATLAS Computing For Today's Networks *J. Phys.: Conf. Ser.* **396** 032019

[9] The MONARC Project, http://MONARC.web.cern.ch/MONARC/

[10] Elsing M, Goossens L, Nairz A and Negri G 2010 *J. Phys.: Conf. Ser.* **219** 072011

[11] Maeno T et al. (for the ATLAS Collaboration) 2012 Overview of ATLAS PanDA Workload Management *J. Phys.: Conf. Ser.* **331** 072024

[12] Branco M, Cameron D, Gaidioz B, Garonne V, Koblitz B, Lassnig M, Rocha R, Salgado P and Wenaus T 2008 Managing ATLAS data on a petabyte-scale with DQ2 *J. Phys.: Conf. Ser.* **119** 062017

[13] Stewart G A et al. (for the ATLAS Collaboration) 2012 Advances in service and operations for ATLAS data management *J. Phys.: Conf. Ser.* **368** 012005

[14] Schovancová J (for the ATLAS Collaboration) 2012 ATLAS Distributed Computing Monitoring tools after full 2 years of LHC data taking *J. Phys.: Conf. Ser.* **396** 032095

[15] Garonne V et al. 2012 The ATLAS Distributed Data Management Project: Past and Future *J. Phys.: Conf. Ser.* **396** 032045

[16] Maeno T et al. (for the ATLAS Collaboration) 2012 Evolution of the ATLAS PanDA Production and Distributed Analysis System *J. Phys.: Conf. Ser.* **396** 032071
[17] Nowak A 2013 Opportunities and choice in a new vector era *(these proceedings)*
[18] Kama S 2013 Optimizing the ATLAS code with different profilers *(these proceedings)*
[19] Van Gemmeren P, Binet S, Calafiura P, Lavrijsen W, Malon D and Tsulaia V (for the ATLAS Collaboration) 2012 I/O Strategies for Multicore Processing in ATLAS *J. Phys.: Conf. Ser.* **396** 022054
[20] Crooks D et al. (for the ATLAS Collaboration) Multi-core job submission and grid resource scheduling for ATLAS AthenaMP *J. Phys.: Conf. Ser.* **396** 032115
[21] Barreiro Megino F H et al. 2012 Exploiting Virtualization and Cloud Computing in ATLAS *J. Phys.: Conf. Ser.* **396** 032011
[22] Bird I 2013 Grid AND Cloud evolution from now onward *(these proceedings)*
[23] Bauerdick L et al. 2012 Using Xrootd to Federate Regional Storage *J. Phys.: Conf. Ser.* **396** 042009