

Added value of the new features of the ATLAS Computing Model and a shared Tier2&Tier3 facilities from the Community point of view

Gabriel Amorós Vicente

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: Gabriel.Amoros@ific.uv.es

Santiago González de la Hoz

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: Santiago.Gonzalez@ific.uv.es

Álvaro Fernández Casaní

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: Alvaro.Fernandez@ific.uv.es

Miguel Villaplana Pérez

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: Miguel.Villaplana@ific.uv.es

José F. Salt Cairols

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: salt@ific.uv.es

Mohammed Kaci

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: kaci@ific.uv.es

Alejandro Lamas

IFIC (CSIC/UV)

Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain

E-mail: Gabriel.Amoros@ific.uv.es

Elena Oliver*IFIC (CSIC/UV)**Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain**E-mail: Elena.Oliver@ific.uv.es***Javier Sánchez***IFIC (CSIC/UV)**Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain**E-mail: Javier.Sanchez@ific.uv.es***Victoria Sánchez***IFIC (CSIC/UV)**Edificio Institutos de Investigación – 22085, E-46071 Valencia, Spain**E-mail: Victoria.Sanchez@ific.uv.es*

The ATLAS computing and data models have moved/are moving away from the strict MONARC model (hierarchy) to a mesh model. Evolution of computing models also requires evolution of network infrastructure to enable any Tier2 and Tier3 to easily connect to any Tier1 or Tier2. In this way some changing of the data model are required. Data taking in ATLAS has been going on for more than one year. The Tier2 and Tier3 facility setup, how do we get the data, how do we enable at the same time Grid and local data access will be presented in the EMI/gLite middleware flavour. And explained from the point of view of the user community, which is supported at the national and international level also with different Virtual Organizations in addition to ATLAS. Finally, an example of how a real physics analysis is working at these sites will be shown.

*EGI Community Forum 2012**Leibniz Supercomputing Centre (LRZ), Munich, Germany**March 26 – March 30, 2012*

1. Introduction

The *Instituto de Física Corpuscular* (IFIC) is a High Energy Physics Institute where several groups use the same infrastructure to foster their calculations. Besides that, the different fields of the groups translate to different requirements, which can be covered depending of the capacity of the group. For instance some groups require parallel computing, another high storage capacity, others need a very good communication network and other need a fast computing facility. This heterogeneity is very common in centres with a large number of groups as at the IFIC.

Among the different projects to provide resources, ATLAS is by far the biggest one. IFIC participates in this experiment as one of the members of the Spanish Tier-2 together with *Universidad Autónoma de Madrid* (UAM) and *Institut de Física d'Altes Energies* (IFAE) in Barcelona. ATLAS is one of the experiments that have been working within the program of LHC (Large Hadron Collider) [1] at CERN [2] in Geneva. LHC runs protons beams in opposite directions in order to provide enough energy to saturate the energy limits where the Higgs particle in case it exists (the key particle that makes the particle to have weight), can be produced and, additionally, to look on the limits of the Standard Model of particles. The proton beams collide in ATLAS with a total collision energy of 7 TeV. The integrated luminosity recorded at the end of 2011 was 5.25 fb^{-1} . ATLAS has exported more than 25k TB of data to Tier-1s since January 2010. This volume of data is one of the main reasons for the development and establishment of a computing model based on Grid technologies [3] at the LHC experiments. The model handles everything from the storage of raw events at CERN to physics analysis on refined data at home institutes. A proper sizing and organization of the resources, an efficient mechanism to access the data or robust algorithm development are examples of key items in the so-called solution for the steady-state period during the data taking. However, the solution is a dynamic one, and at the times where it was proposed there was not example of any kind similar. We have already several years of experience of simulation and real data, and some lessons have been learned. This evolution affects the way the high energy users, and in particular the ATLAS users, understand the infrastructure and process their analysis. We present how this evolution affects to the IFIC users, particularly to the ATLAS users. We present how IFIC infrastructure is organized, how is working and how it has been affected by the new changes in the ATLAS computing model. At the last section we present an example of physics analysis in order to show the modifications respect to the previous analysis workflow.

1.1 IFIC Spanish Tier-2

The Worldwide LHC Computing Grid project (WLCG) groups the different types of computing centres in a tiered hierarchy that ranges from the Tier-0 at CERN to the 11 Tier-1 centres and the 70 Tier-2 centres distributed world wide. A Tier-1 centre, which leads its associated Tier-2s, forms with them, an ATLAS cloud. The roles of the different tiers are well established in order to produce a stable and efficient distributed computing facility for the ATLAS experiment [4]. This distribution in Tiers is a logical one because of the central point of CERN as site of the LHC; simulated and real data are obtained from the Tier-0. The data is

accumulated and have to be stored for years, implying a list of Tier-1s as a backup system besides additional general data processing activities. The final users, the physicists, manage the data on the Tier-2s for their analysis, requesting any needed data. In general the Tier-2s are unique centres, however the features of the ATLAS community in Spain, made more efficient to create a federated Tier2 with three institutions: IFAE, UAM and IFIC, reaching around 5% of the total ATLAS resources [3]. IFIC represents 50% of the ATLAS Spanish resources and has the responsibility to coordinate the activities of the Spanish Tier-2 federation. In table 1 the evolution of IFIC Tier2 resources is shown.

Table1: Pledged evolution of IFIC Tier-2 resources. The numbers in 2011 (bold) correspond to the current deployed resources.

Year	2010	2011	2012	2013
CPU(HS06)	6000	6950	6650	7223
DISK(TB)	500	940	1175	1325

2. Evolution of the Atlas data and computing model

The previous ATLAS model was organized hierarchically; although the T0 and T1 could communicate with each other, that was not the case for the T2; they could communicate directly only with the own T1, the own cloud; in the case of Spain the ATLAS T1 is located in Barcelona at the *Port de Informació Científica* (PIC). The consequence of that is that any request of a dataset not located at the own T2 but in the cloud of a T2 belonging to another T1 cloud, have to be transferred to the own T1, then to the far T1 and then to the far T2. This is a good vision in the ideal case of good communications and resources. The experience shows that this is not the case and the real world imposes a modification on the computing model. With the new one, Tier-2s with a good network connection (T2Ds) are allowed to connect to other Tier-1 sites and to Tier-2 sites from a different cloud. The CPU resources are more efficiently used in this way and high priority tasks can be done more quickly. IFIC Tier-2 is considered to be a T2Ds and is now working as a multi-cloud site (Tier-2s with a good network connection are allowed to connect to other Tier-1s and to Tier-2s from a different cloud). Therefore, there are direct transfers from/to all ATLAS Tier-1.

The revised data distribution system is composed of pre-defined distribution based on the model, dynamic data placement based on the usage, and on-demand replication. The pre-defined data distribution creates “primary” replicas (to be available on disk) at Tier-1 sites for redundancy and at Tier-2 sites for end-user analysis, as well as “secondary” replicas (they are extra replicas that are created for supposedly popular data using the remaining available disk spaces) to give larger opportunity for analysis. The replicas at the Tier-1 sites are either exported from the Tier-0 or copied from the other Tier-1 sites. The replicas at Tier-2s are created from the replicas at the Tier-1 sites. The dynamic data placement is implemented in the distributed analysis system to increase “secondary” replicas based on the usage of the data, in order to reduce the waiting time of analysis jobs

2.1 Flattening the model

The original model was a working starting point but Tier-2 activity was strongly linked to the associated Tier-1 reliability. Some Tier-1s did not have associated Tier-2s and had few tasks to process while they had the storage to host a significant fraction of task outputs. In addition, Tier-2 had the computing resources to do reprocessing but was limited due to required direct access to the Tier-1 database. As a consequence, some of the sites were not used at full capacity, especially Tier-2s. Some of the tiers were not fulfilling the required level of utilization, especially Tier-2 like the one at IFIC. This situation was mainly due to two issues, being one of them operational and the other related the fact that the data flows among the Tier-2s were limited [5].

On the operational level, one question that limited the usability was the dependency on Tier-1 in order to run jobs. As a consequence, when our Tier-1 was in scheduled downtime the associated Tier-2 were affected.

On the other hand, job input files, as well as IFIC job outputs had to be transferred to associated Tier-1. For example, if some data produced at IFIC was needed in a Tier-2 at another cloud, dataset had to be transferred via source and destination associated Tier-1. This situation affects the use of Grid by the users because, due to the size of the datasets that could need, they had to be sure these datasets are in the sites where their analysis will be running. This is a task that makes the Grid not transparent to the user, and then forcing them to learn about data management tools out of their research field.

The solution to improve this situation was taken by ATLAS in the form of flattening the model from a tier to a mesh. Now Tier-2 can directly exchange data with Tier-1 of different clouds, and even with other Tier-2 globally. However these relations are scored considering the real situation of the communications and resources. This situation is obtained with tests related to the crucial aspects: connectivity and availability. Periodically tests of the connections of each pair Tier-2/1-Tier-2/1 are performed in order to realize about the real state.

2.2 Connectivity

Network is a key component in the evolution of the ATLAS model for the Tier-2 [4], as they have to be well connected to be able to exchange data. In order to test connectivity, ATLAS monitors transfers among sites.

In terms of connectivity, the concept of so-called T2D was introduced. In three months, the connectivity of each Tier-2 with respect to each one of the Tier-1s is reviewed. Those that fulfil a certain performance are then classified as T2D. Only the T2D will be used to run jobs that were attributed to a cloud different from the one they belong to. At the moment, the limit is set to an overall transfer performance of at least 5MB/s to at least 80% of Tier1s in ATLAS, for very large files (>1GB).

IFIC has fulfilled transfer performance requirements as in shown in Figure 1, even though fluctuating values for bigger files and transfer to remote sites have been observed. For this reason, we have installed new GridFTP servers. As the transfers are now distributed among more servers IFIC is getting more stable values. Thank to this, IFIC obtained a T2D status. The tests were performed with different sizes of files to provide file independent state of the transfers.

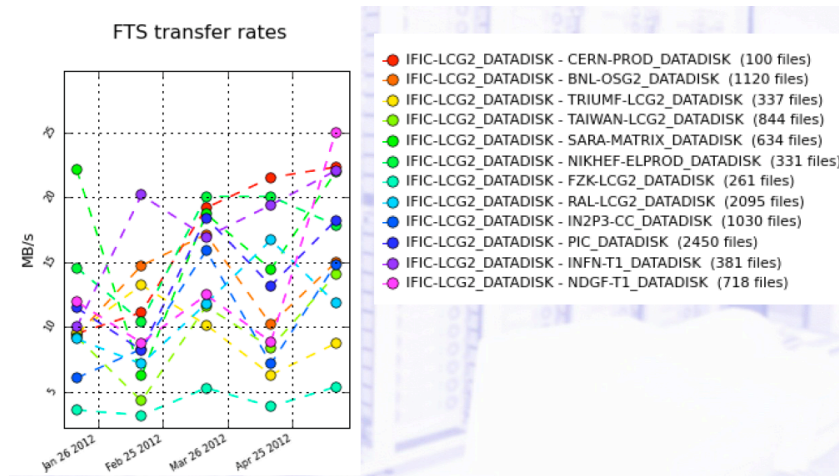


Figure 1: T2D performance from IFIC to Tier-1s.

2.3 Availability

The other important point for a Tier-2 site is high availability. ATLAS uses the HammerCloud framework to test site availability, by constantly submitting typical analysis jobs to every site. Site availability for analysis jobs is defined based on the “analysis panda queue status”: (the amount of time queue status = online in the period) / (the whole period).

Sites are evaluated monthly and ranked in 4 groups (Alpha, Bravo, Charlie and Delta) depending on its availability and its connectivity:

alpha – availability >90% if site is also a T2D (~30 sites)

bravo - availability >90% but site is not a T2D (~20 sites)

charlie - availability >80% (~5 sites)

delta - availability <80% (~10 sites)

The indicated numbers of sites correspond to the status of March 2012.

Scheduled downtimes of a site and unavailability not due to the site are not considered for the classification.

At the moment, around half of the Tier2-s are classified as T2D, and practically all of those as *alpha* as well. Summing up the share of all T2Ds, they obtain at the moment (March 2012) a total share of around 75% of the data distribution [5].

During last months (January-June 2012) IFIC has been qualifying as Alpha Site and has been available more than 90% of the time. Additionally the IFIC resources are monitored with the EGI (European Grid Infrastructure) tools; these tools check the situation of the different VOs. However, the highly demanding ATLAS experiment forces a response time smaller than other VOs. This situation is related to the very collaborative and mature implementation of the tools for data management and processing within ATLAS collaboration. Besides that, the pledges for the institutions with ATLAS are tied in resources but also with support, implying a fast response. This is a common frame in high energy physics experiments that explains the early acceptance and adoption of a collaborative Grid by the physicists, that are getting more and more used to the tools and the environment to perform the work. On the other hand other VOs are getting more and more experienced in the use of the collaborative grid, also benefiting

from this framework and the involved processes, which eventually improves the experience for their users.

2.4 Dataset replication

Being an Alpha T2D site, IFIC is now able to replicate higher amounts of datasets. As can be seen in Figure 2, the Spanish Tier-2 sites (ES-Cloud) are getting now more datasets than Tier-1 (PIC) while, in the past, it used to be the Tier-1 the one getting more datasets. Now the users can have more chances to have the datasets close to the site where their analysis will be performed. The data transfers are not tied to the cycle T2-T1-T1'-T2', and the possibility of a direct transfer T2-T2' decreases the waiting time for availability of the datasets at the target site.




Figure 2: Data transfer volume in the Spanish Tier-1 and Tier-2 sites during January and February 2011(top) and 2012(bottom).

The same behaviour is observed in other Tier-2s for the ATLAS Collaboration. Tier-2s now get more datasets than Tier-1s because the disk size in Tier-2s has increased significantly while Tier-1 total size was not increasing as much [5].

3. IFIC Infrastructure and operations

Here we show a short summary of the resources and achievements at IFIC. The fact that IFIC is well qualified, in connectivity and availability as was shown in figure 1, implies that the connections and resources have enough quality. On the other hand the management and support for these resources is important as well.

3.1 Computing resources

IFIC computing resources are distributed in two areas: the resources devoted exclusively to the ATLAS, and the ones provided within the GRID-CSIC project by the Spanish Research Council. The latter resources can be used for a variety of projects, and currently the usage follows the distribution of 25% for IFIC users, 25 % for CSIC projects, 25% for EGI Virtual Organizations, and 25% for Iberian Grid virtual organizations, that are also included in EGI. The GRID-CSIC resources are separated from the ATLAS resources; the second ones are dedicated due to the pledges of the experiment. However the GRID-CSIC resources are general ones and can be used by the VOs supported by IFIC, and for instance, it can happens that ATLAS users at IFIC could use GRID-CSIC resources. The figures are shown in table 2.

Table 2: IFIC resources.

	IFIC ATLAS Tier - 2	IFIC GRID - CSIC
Cores	640 ¹	1232
Storage	940 TB	180 TB

3.2 Storage resources

Last year we hit the Petabyte of storage for ATLAS, plus 180 TB more for the rest of the supported VOs. We are using Lustre v1.8 [6] as a backend posix file system and with this release we included pool capabilities to the installation. This allows us to partition the hardware inside a given file system in order to have a better data management. In addition to this, we can assign selected Object Storage Targets (OST) to an application/group of users, and we can separate heterogeneous disks in the future. We have four file systems devoted to different Grid virtual organizations:

- (1) */lustre/ific.uv.es*. For the Virtual Organizations. Mounted *read only* on worker nodes (WN) and user interfaces (UI), and *read/write* on GridFTP + SRM.
- (2) */lustre/ific.uv.es/sw*. Devoted to software, mounted *read/write* on WNs and UIs. This is what ATLAS used before migrating to CVMFS system.
- (3) */lustre/ific.uv.es/Grid/atlas/t3* Space for T3 users. *Read/write* on WNs and UIs.
- (4) *xxx.ific.uv.es@tcp:/homefs* on */rhome* type lustre. Shared home for users and MPI applications. *Read/write* on WNs and UIs.

Users in general will access the storage resources with (1). (2) Is for software updates, installations and with administrator roles. (4) Specific for MPI. (3) Is for a local Tier-3 facility

¹ ATLAS uses the normalization to HEP-SPEC06. In this case 640 cores = 6950 HEP-SPEC06. See Table 1.

to help the local users analysis. This separation provides a good flexibility for the data management.

Storm is used as a SRM server for the Grid access. We developed an authorization plugin in order to respect local file system with the corresponding user mappings and ACL's. Lately, we have seen performance degradation due to the increasing number of connections. The storm server has been upgraded to 8-core machine. However, we still see this service as a point of failure in many incidents, and we are working to find a better solution in the future.

3.3 Summary of utilization

In 2011, we supported 22 Virtual Organizations including ATLAS, our local institute virtual organization IFIC, European projects like vo.agata.org, and t2k.org, and all the virtual organizations of the IberGrid Federation. More than 3.5 million jobs have been executed, 6 million CPU hours consumed, corresponding to close to 14 million KSi2K normalized CPU time. More than 90% of these resources are accounted to ATLAS, which runs on our resources all year long.

4. Example of a high energy physics analysis in Grid

Until now we have shown the situation and the changes on the infrastructure of the new computing model for ATLAS. In fact we have talked about the new computing model but we did not explain how it affects to the users. From the point of view of the infrastructure the main difference is a better monitoring of the resources and the special tests between pairs of T1/2. For the rest, the pledges remain the same, that is, there is not modification with the required computing and storage resources; it is up to the institution to be promoted to a better ranked one in the relations with other T1/2; the reality of the communications and resources imposes a better use of them, and then a better efficiency.

Now we change the perspective from an infrastructure and service point of view to the end-user point of view. In our case, the end-users are the physicists working daily on physics analysis. Typically, a physics analysis has two parts. In the first stage, physicists run an analysis program that uses a given number of collision events. These events can be stored in different datasets that are usually spread at different sites. At this step, the Distributed Computing and Data Management Tools, based on Grid Technologies, are used in an exhaustive way. The output of this first step is often a set of ROOT [7] ntuples or properties vectors.

In the second stage, the physicists analyse the ntuples interactively in order to get the final plots, to refine the analysis, etc. In our institute the physicists can also take advantage of using both Tier-2 and Tier-3 infrastructures. Initially the Tier-3s infrastructure was not foreseen in the ATLAS computing model, however the requirements of the users of powerful resources for their own and local analysis tasks made a reconsideration of Tier-3. Some words are worthy to explain the Tier-3.

As an extra service outside the backbone of the ATLAS computing model, the implementation of this new Tier is a decision of the institute and then the variety of Tier-3 depends on the needs of the centres. At IFIC the Tier-3 resources are split into 2 parts (see ref. [8, 9]):

- Resources coupled to IFIC Tier-2, which are included in the Grid environment and used by the IFIC-ATLAS users. Only if the resources are idle, the whole ATLAS community can use them.
- A computer farm to perform interactive analysis (PROOF [10]), which is outside the Grid Framework.

It is remarkable that, following the criteria of efficiency of use of the resources, resources with less use can be added to the Tier-2; to avoid the minimal disturbance on the process of connect/reconnect, the lustre servers for Tier-3 are closely located to the Tier-2 ones. A schematic view of this Tier-3 model can be seen in Figure 3.

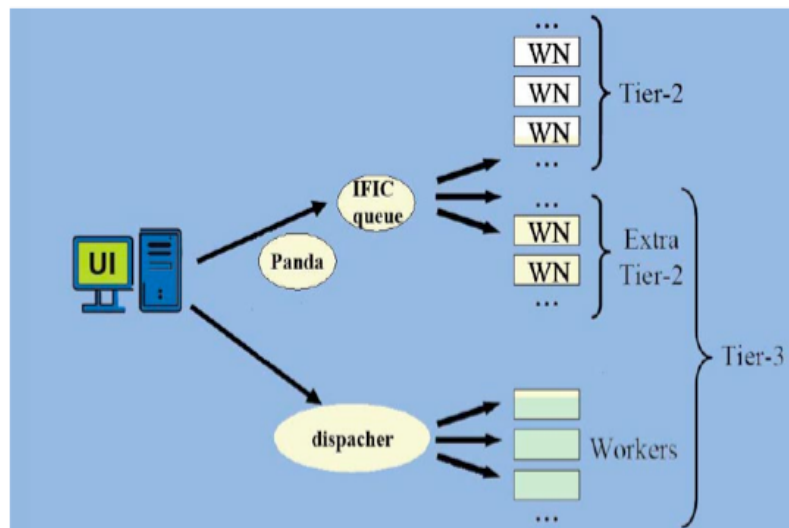


Figure 3: Schema of the IFIC Tier-3

Resuming the Distributed Analysis process, in general is necessary a use of the following ATLAS tools:

1) For Data Management:

- Don Quijote 2 (DQ2) [11]: to obtain information about data and to download and register files on Grid.
- DaTri: the end-user dataset subscription service.
- AMI (ATLAS Metadata Interface)

2) For GRID Jobs:

- PanDA Client [12]: analysis job submission tool.
- GANGA [13] (Gaudi/Athena and Grid Alliance): it is a job management tool for local, batch system and the Grid.

The workflow is the following: the user needs a specific number of events in several datasets; to have them on the computing resources running their analysis, they look for the datasets with AMI and/or DQ2, later they ask for them officially with DaTri and then copy locally to their own institute, and DQ2 can be used to manage the datasets on Grid. This is the standard way of proceeding, however the new computing model can save some of these steps.

We can go deeper in the daily user activity in distributed analysis by illustrating it with an example [9]. Simulated and real data input files represent a volume of several TB of information. On the other hand, the analysis activity workflow can be divided as follows:

- 1) Test the analysis locally, where the input files are downloaded with DQ2 tools.
- 2) Submit a first job to Grid which will create an output file with reduced information, the input can be either real or simulated data and the run time for a typical job is around 20 hours;
- 3) Submit a second job to Grid (normally at the Tier-3) with the objective of doing a refined analysis (reconstruction, application of cuts and selections, etc.). In this case the input is the output of the first job (phase 2) and the execution time is around 2 hours.

If we focus on phase 2, the schematic view of the activity can be seen in Figure 4. Working within the GANGA Framework, a python script is created in order to give the requirements of the job, characteristics such as application location, input-output, a replica request to IFIC or the output files splitting. The script allows sending the job to the Grid and, once the jobs finish successfully, the output files are copied to the IFIC Tier-3.

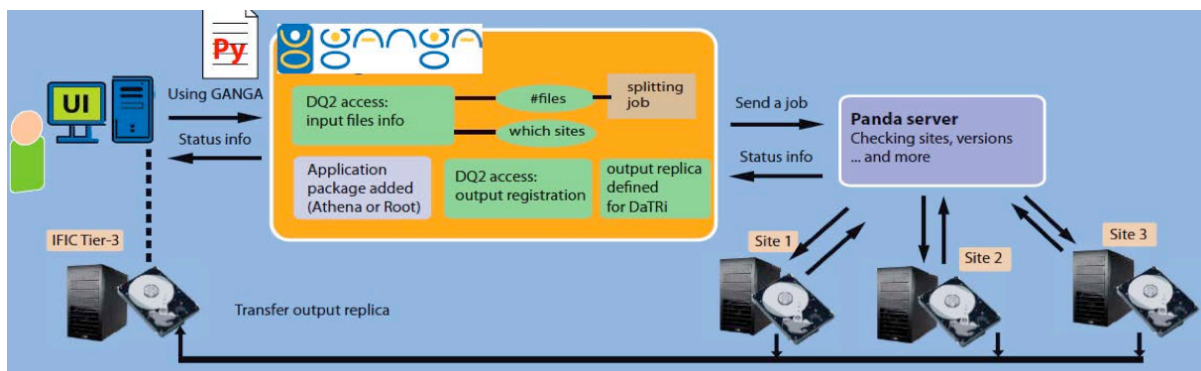


Figure 4: Schematic view of the Distributed Analysis activity (phase 2)

This typical workflow is the same in both the previous and the new computing model. However within the new one the analysis tools have been improved; the user just has to select the datasets needed without worrying about the existence of a copy within the sites where to run the analysis. With the previous model the steps of looking for the sites with better resources to run the analysis was restricted to the sites where there exist copies of the datasets. With this behaviour the real world is posed to the ideal one of ideal resources. This is against the ideal image of the Grid as an infrastructure where the user does not have to know about the state of the infrastructure layer. With the model there have been some steps on the direction of this image because, although it remains the selection of resources with good rates of execution time of the Grid jobs, the task of selecting the resources following the datasets existence is over passed. This is a good advance in the direction of helping the users to do their job without forcing them to be experts in Grid.

The revised data distribution system is composed of pre-defined distribution based on the model, dynamic data placement based on the usage, and on-demand replication. The pre-defined data distribution creates “primary” replicas (to be available on disk) at Tier-1 sites for redundancy and at Tier-2 sites for end-user analysis, as well as “secondary” replicas (they are

extra replicas that are created for supposedly popular data using the remaining available disk spaces) to give larger opportunity for analysis. The replicas at the Tier-1 sites are either exported from the Tier-0 or copied from the other Tier-1 sites. The replicas at Tier-2s are created from the replicas at the Tier-1 sites. The dynamic data placement is implemented in the distributed analysis system to increase “secondary” replicas based on the usage of the data, in order to reduce the waiting time of analysis jobs.

For a heavy exotic particles analysis at IFIC, in two weeks 6 users sent 35728 jobs with success to 64 sites (Tier-2s and Tier-1s). 1032 jobs were sent to the Spanish Tier-2 (2.89%) and 815 datasets were used as input files producing 1270 files as output.

From the point of view of the user, functionally of dataset replication is transparent to users, but eventually they learn that it is in action because:

- A job was finally sent to a destination that did not originally had the requested datasets;
- Later using dq2 tool the original datasets were copied with the dynamic data placement algorithm.

5. Conclusion

IFIC had adapted well to the changes in the computing model. Thanks to its excellent connectivity and availability, IFIC obtained the Alpha T2D status. This new status implies an increase in the number of jobs and data transfers where IFIC infrastructure performs well. Besides official production tasks, the new status facilitate the transfers of datasets directly from others T1/2, allowing bigger flexibility for the users. Other ATLAS Tier-2s have observed similar results. Tier-2 activities are now less dependent to Tier-1 and participate to more critical activities in ATLAS. Tier-2s are again receiving data immediately to have higher contribution to analysis activities. Their availability is reported every month, requiring a good connectivity for data transfer and production.

From the user point of view the improvements of the analysis tools hides some of the tasks of data management on the Grid helping the physics community to focus on their analysis tasks.

A real analysis example has been discussed where, apart from the distributed analysis, an interactive analysis runs at the local Tier-3. This example is the best proof of the good throughput of the ATLAS Computing Model and the Tier-2 and Tier-3 at IFIC. It is also remarkable that the users of other VOs are not affected on these changes.

Acknowledgements

We acknowledge the support of MICINN, Spain (Plan Nacional de Física de Partículas FPA2010-21919-C03-01)

References

- [1] ATLAS Collaboration, G. Aad et. al., The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003.
- [2] CERN, European Organization for Nuclear Research, <http://www.cern.ch>
- [3] “The ATLAS Computing Model”, D. Adams et al., ATL-SOFT-2004-007, CERN, (2004)
- [4] “ATLAS Distributed Computing Operations in the First Two Years of Data Taking” Ueda, I for the ATLAS Collaboration, ATL-COM-SOFT-2012-007.- Geneva : CERN, 2012 - 9 p.
- [5] “The evolving role of Tier2s in ATLAS with the new Computing and Data Distribution model”, ATL-SOFT-PROC-2012-010.- Geneva: CERN, 2012 9 p
- [6] Lustre File System, <http://wiki.lustre.org>
- [7] ROOT, <http://root.cern.ch/drupal>
- [8] S. Gonzalez de la Hoz et al. “Analysis facility infrastructure (Tier-3) for ATLAS experiment”, Published in Eur.Phys.J.C54:691-697, 2008.
- [9] M. Villaplana et al. “First tests with Tier-3 facility for the ATLAS experiment at IFIC (Valencia)”, ISBN 978-84-9745-549-7, pages 212-220.
- [10] M. Ballintijn et al. “The PROOF Distributed Parallel Analysis Framework based on ROOT”, arXiv.org: physics/0306110, 2003
- [11] “Managing ATLAS data on a petabyte-scale with DQ2”, M Branco et al., J. Phys.: Conf. Ser., 119 062017, 2008
- [12] “Proceedings of XII Advanced Computing and Analysis Techniques in Physics Research”, P. Nilsson et al. Proceedings of Science, 2008
- [13] “Ganga: a tool for computational-task management and easy access to Grid resources”, F. Brochu et al. CoRR, abs/0902.2685, 2009