The ATLAS Data Challenges

G. Duckeck, O. Biebel, T. Christiansen, J. Dubbert, J. Elmsheuser, M. Erlebach, F. Fiedler, R. Hertenberger, J. Kennedy, R. Mameghani, D. Merkl, T. Nunnemann, M. Obermaier, F. Rauscher, D. Schaile, M. Schott, A. Staude, H. Steffens, R. Ströhmer, A. Varga, F. Vollmer, and C. Zupancic

The LHC experiments will require huge computing resources in order to store, reconstruct and analyse the data recorded. The workload will be distributed to many regional computing centers in the participating countries. The automatic management of these data sets and the effective workload distribution pose non-trivial problems. To aid with the development of this distributed computing infrastructure a series of so-called data challenges (DC) started in 2002. These test productions allow us to evaluate the Computing Model, the ATLAS software suite, the Data Model and also to ensure the correctness of technical choices made for ATLAS computing. In 2004 the second large-scale test production started with one of its major aims being the use of GRID tools to their maximum extent and the assessment of their ability to process events sets on a LHC scale. The goal was to fully simulate 10^7 physics events from various processes by using computational resources at about 100 institutes worldwide.

In Germany two centers participated: The GridKa center in Karlsruhe and the Leibniz Rechen-Zentrum (LRZ) in Munich. GridKa is a dedicated center to serve as Tier-1 facility for the four LHC experiments [1]. It was established in 2002, it took part already in previous data challenges and supports also four non-LHC experiments (BaBar, CDF, Compass and DØ). GridKa is part of the LCG production Grid [2] and thereby the prerequisites for DC2 production are fulfilled.

The LRZ facility in Munich is a candidate for an ATLAS Tier-2 site, operated jointly by LMU, LRZ, MPI and RZG. The participation in DC2 production serves as a pilot project involving members of LMU Munich and LRZ. At present the integration of the LRZ cluster into the LCG production Grid causes problems because the current version of the LCG Grid middleware packages impose strong requirements on the cluster setup. These can hardly be fulfilled by existing sites which need to continue the support of other applications. As an alternative we have deployed at LRZ the NorduGrid/ARC Grid middleware [3], which is one of the three Grid flavors supported by ATLAS. Figure 2 shows the web-based monitoring tool for the NorduGrid resources. The ARC middleware has been adapted for use on a large variety of Linux flavors and batch systems. A particular advantage of ARC is the transparent installation on a single front-node server only.

For the production in DC2 a dedicated production system has been developed which splits the full dataset is split into several 100,000 jobs. It is described in more detail in [4]. The execution of these jobs in NorduGrid proceeds as follows:

- The NorduGrid executor contacts a central database to retrieve the information needed for the next bunch of jobs (typically 1000).
- The jobs are distributed to the sites registered in NorduGrid. The ARC resource broker takes care of a proper distribution of jobs to sites which match the requirements in terms of CPU, memory, etc.
- Before a job starts on the local sites any input files needed are downloaded using Grid copy-tools as specified in the job description.
- Jobs are submitted to the local batch system. During execution access to remote databases is needed.
- After the job has finished the data files produced are uploaded to a central NorduGrid storage site. The files are registered in a Grid replica catalog and made available worldwide for further processing steps.

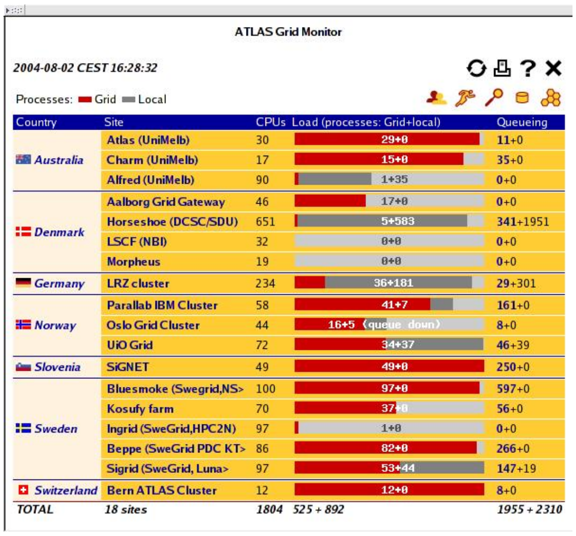
ATLAS Grid Monitor

2004-08-02 CEST 16:28:32

Processes: Grid Local

Country	Site	CPUs	Load (processes: Grid+local)	Queueing
Australia	Atlas (UniMelb)	30	29+0	11+0
	Charm (UniMelb)	17	15+0	35+0
	Alfred (UniMelb)	90	1+35	0+0
Denmark	Aalborg Grid Gateway	46	17+0	0+0
	Horseshoe (DCSC/SDU)	651	5+583	341+1951
	LSCF (NBI)	32	0+0	0+0
	Morpheus	19	0+0	0+0
Germany	LRZ cluster	234	36+181	29+301
Norway	Parallab IBM Cluster	58	41+7	161+0
	Oslo Grid Cluster	44	16+5 (queue down)	8+0
	UiO Grid	72	34+37	46+39
Slovenia	SiGNET	49	49+0	250+0
Sweden	Bluesmoke (Swegrid,NS>	100	97+0	597+0
	Kosufy farm	70	37+0	56+0
	Ingrid (SweGrid,HPC2N)	97	1+0	0+0
	Beppe (SweGrid PDC KT>	86	82+0	266+0
	Sigrid (SweGrid, Luna>	97	53+44	147+19
Switzerland	Bern ATLAS Cluster	12	12+0	8+0
TOTAL	18 sites	1804	525 + 892	1955 + 2310

Fig. 1: Snapshot of ATLAS test production in NorduGrid

The results of the actual DC2 production were mixed: Many fundamental changes had been made to the experiment specific software, e.g. the detector description, the data management and the full transition from Fortran to C++. These have been tested for the first time on a large scale in DC2 and proved to be very successful. We also managed to produce the planned 10^7 events entirely using Grid tools. However, there also had been substantial problems, e.g. the failure rate was rather high, the overall throughput did not meet the original expectations and several bottlenecks limited the production. A major further effort is needed in the next years in order to have a fully functional system available in time for the start of LHC.

The participation of the German sites was satisfactory, about 1200 jobs were executed at LRZ and about 25000 at GridKa. We gained important experiences for more continuous and larger scale productions in the future.

References

[1] Requirements for a Regional Data and Computing Centre in Germany, `http://grid.fzk.de/grid/LHCComputing-1july01.pdf` (2001)
[2] LHC Computing Grid Project, `http://lcg.web.cern.ch/LCG/`
[3] NorduGrid/ARC, `http://www.nordugrid.org/`
[4] J. Kennedy *et al.*, this report, p.97