Slovak Participation in the World LHC Computing Grid

Jan Astalos[1], Marian Babik[1,3], Robert Breier[2], Dusan Bruncko[3], Peter Chochula[2], Vladimir Fekete[2], Ingrid Kulkova[3], Ladislav Sandor[3], Branislav Sitar[2], Pavel Stavina[2†], Miloslav Straka[3], Pavol Strizenec[3], Stanislav Tokar[2], Martin Vala[3], Jana Vrlakova[4], Tibor Zenis[2]

[1] Institute of Informatics SAS, Bratislava, Slovak Republic
[2] Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovak Republic
[3] Institute of Experimental Physics SAS, Kosice, Slovak Republic
[4] Faculty of Science, Pavol Jozef Safarik University, Kosice

Abstract. Slovak Republic participates on two experiments on the Large Hadron Collider (LHC) ATLAS and ALICE. Slovakia participates also in the Worldwide LHC Computing Grid (WLCG). The article presents the current situation in Slovak involvement in WLCG project and gives some perspective for the near future.

1 Introduction

The Large Hadron Collider (LHC) is a world biggest particle accelerator, built at the European Organization for Nuclear Research (CERN) [1]. Providing the particle beams with highest energy, achieved till now in laboratory, it serves the physics community in their searches for basic properties of matter. The machine is installed in circular tunnel with a circumference of 27 km, with two particle beams circulating in opposite directions (Fig. 1), with target energy 7 TeV for proton beam. Until end of 2011 the LHC is running at the half of the target energy - 3.5 TeV per proton beam. The particles in the beams are formed in bunches (there will be 2808 bunches per beam), which will collide every 25 ns. Two beams are intersecting on four places on the ring, where the detectors are installed. These detectors are measuring the products of particle collisions, trying to answer variety of fundamental physical questions. Slovak Republic participates on two of those detectors - ATLAS [2] and ALICE [3] from the beginning of their construction.

Both ATLAS and ALICE detectors (see short description below) are producing enormous amount of data, which needs to be stored, processed and subsequently many times analyzed by many groups of physicists. The Worldwide LHC Computing Grid (WLCG) project [4] has been created to cover computing needs of the LHC experiments.

†deceased

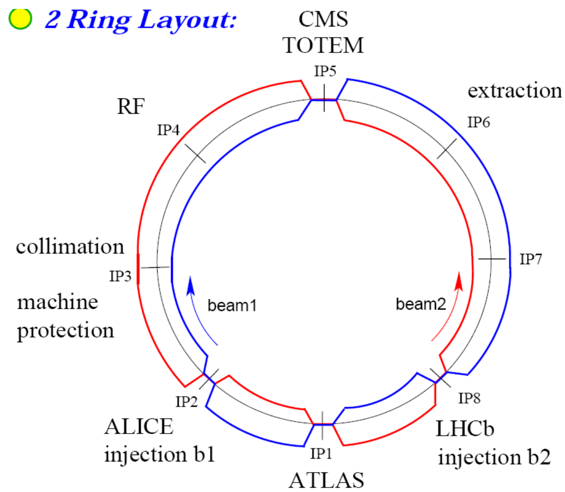

Fig. 1. Schematic view of LHC. Contra rotating beams are crossing in the four interaction points, where experiments ATLAS, ALICE, CMS and LHC-b are located.

1.1 ALICE experiment

The ALICE experiment is dedicated to study the strongly interacting matter under extreme conditions (temperature, energy density), where the new state of matter - quark-gluon plasma (QGP) - is expected to exist. The heavy ions collisions are used as a primary tool to recreate and study such extreme conditions. Detailed study of QGP properties could answer the questions how the Universe looked like few microseconds after the Big Bang.

The ALICE detector is very complex system (Fig. 2) with more then 15 millions of detecting elements, which allows a simultaneous study of different QGP signatures, manifesting themselves in properties of secondary hadrons, muons, electrons and photons, produced in central heavy-ion (Pb-Pb) collisions. The detector will study also the collisions of lighter ions (Ar) as well as proton-proton and proton-ion interactions.

1.2 ATLAS experiment

ATLAS is a general-purpose particle detector, it is so far the biggest particle detector ever built. It is 46 m long, 25 m in diameter and weights 7000 tonnes. Using the proton-proton, as well as heavy-ion collisions, it will study the fundamental questions of particle physics, like origin of masses of the particles, existence of so called super-symmetry, matter-antimatter balance in the early Universe and so on. ATLAS consists of several sub-detectors, used for measuring either particle

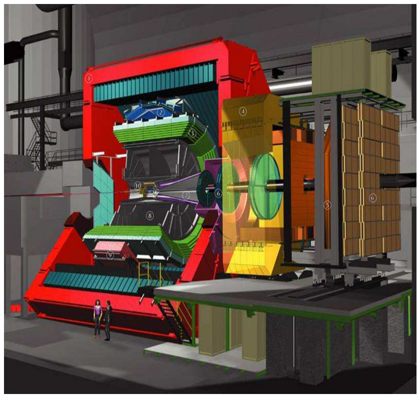

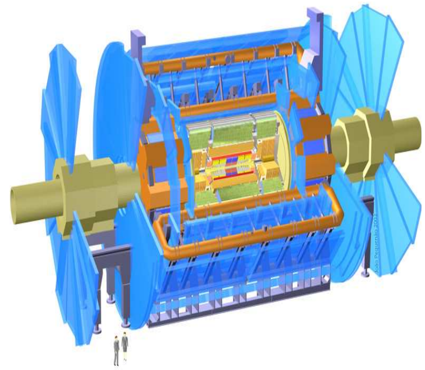

Fig. 2. ALICE schematic view

Fig. 3. ATLAS schematics view

trajectories (Semi-Conductor Tracker, Transition radiation tracker, Muon spectrometer), or particles energy (Liquid-Argon and tile calorimeters). The detector is read-out by more than 160 millions of electronics channels.

1.3 Computing resources requirements

The data flow generated by LHC detectors exceeds by orders of magnitude the amount of data produced by previous generations of high energy physics experiments. The raw data is produced at 40 MHz frequency (which corresponds to collisions rate at LHC). The average event size is 1-2 MB. There is no possibility yet to store such data flow, and also not all of the events are interesting for analysis, therefore a multi-level trigger is used to pick-up the interesting events. On the first level, it's fast and crude algorithm, built in electronics hardware, which choose nearly online the candidate events. The candidates are then checked by more sophisticated algorithms (running on farm of commodity PC's) on next levels (one in case of ALICE, two levels in case of ATLAS).

After trigger filters, the events are written to permanent storage with frequency 100 Hz for ALICE and 140Hz for ATLAS. The total amount of stored statistics is approximately 10^9 events per year in each experiment. All the raw data needs to be subsequently reconstructed, calibrated (one or few times) and analyzed (many times). Complex simulation programs are used to simulate all details of studied processes. This is a demanding task with impact on additional storage and processing capacity. The total CPU need is estimated to app. 100 000 modern CPU's for all LHC experiments. The total data volume is then 6 (11) PB per year for ALICE (ATLAS).

There are also different access patterns to a data, from once store/read for raw data to many read/write for analysis data.

2 Worldwide LHC Computing Grid

In order to meet the needs of the LHC experiments, CERN has initiated the LHC computing grid project [4] back in 2001. The project cooperated with all major grid infrastructures, which were developed at that time, namely Enabling Grid for E-sciensEs (EGEE) [5], Open Science Grid (OSG) [6] and Nordu Grid (NG) [7]

During the years of development, deployment and testing, the infrastructure has proved, that it fulfill the needs of LHC experiments. The architecture is based on hierarchical model of centers (so called Tiers) 2, with top-level Tier-0 located at CERN (concentrated on primary data storage and first pass of data reconstruction), 11 big Tier-1 centers (which are doing most of data re-reconstruction and analysis) and about 38 federations of smaller Tier-2 centers (doing mostly simulation and analysis). There are 35 countries involved in building and running those centers. All those centers are globally monitored, accounted and supported. A number of Tier 3 centers with scale warring from few individual computers up to mid-size farms (tens of CPU's and few TB of storage) is deployed in participating institutes. Those centers provide facilities for physics studies performed by local users and are not centrally monitored.

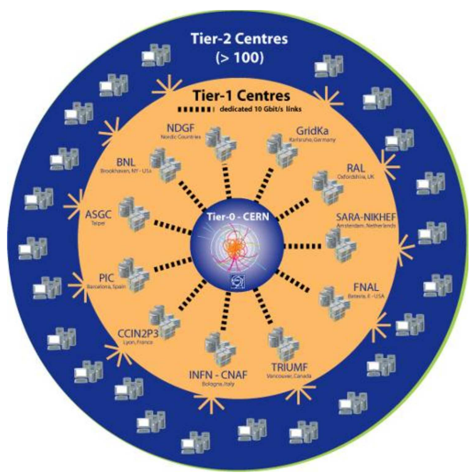

Fig. 4. Schematic view of WLCG global architecture

The schematic view and some basic parameters of central Tier-0 at CERN are shown in Fig. 5 and Tier-1 in Fig. 6

3 Slovak participation in WLCG

There are three research institutions from Slovak Republic actively working on LHC experiments, namely the Faculty of Mathematics, Physics and Informat-

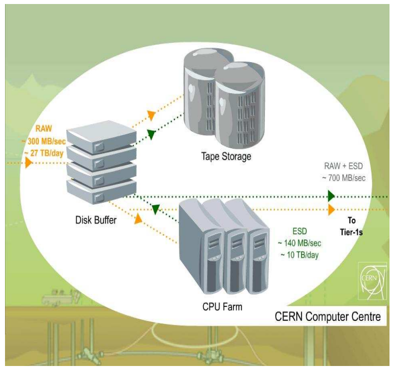

Fig. 5. Schema and parameters of CERN Tier-0

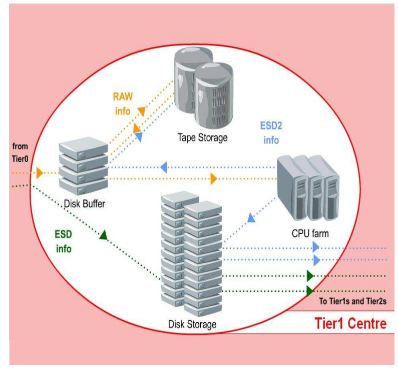

Fig. 6. Schema and parameters of Tier-1

ics, Comenius University, Bratislava (FMFI), Institute of Experimental Physics, Slovak Academy of Sciences, Kosice (IEP-SAS) and Faculty of Science, Pavol Jozef Safarik University, Kosice (UPJS). Participating in computing activities is expected from all institutions collaborating at LHC, and Slovak institutions decided to make their contribution by means of building and maintaining the WLCG grid centers.

There are two clusters in Slovak Republic, both funded by the Ministry of Education within the collaboration with CERN. The clusters are located at FMFI and IEP-SAS (see photos on Fig. 7). The FMFI cluster is supporting only ATLAS and ALICE virtual organizations (VO), IEP-SAS cluster is in addition also accepting jobs from CDF and H1 VO (both are other particle physics experiments in US and Germany, where IEP-SAS is participating).

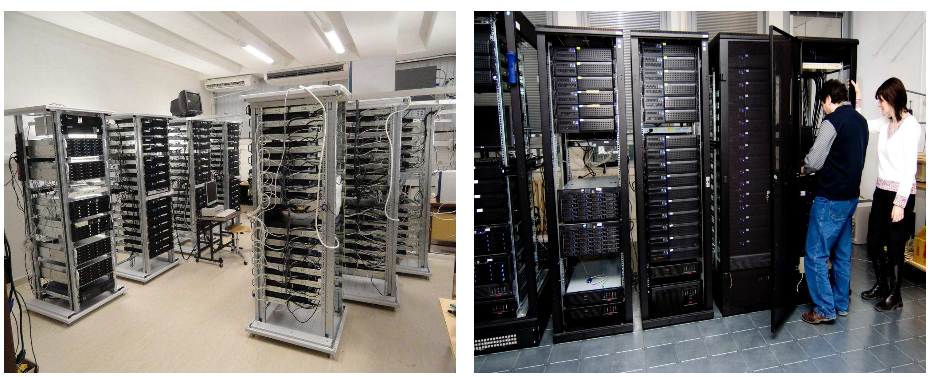

Fig. 7. Clusters at FMFI and IEP-SAS

Slovak clusters are Tier-2/3 (depending in which VO), in production for few years already (centrally monitored [8]) with gradually increased capacity. The current capacity is described in Table 1. The plan is to add more storage and also some CPU, to reach a size of approximately 350 cores per site and over 200 TB disk capacity per site.

Both clusters follow the typical WLCG architecture with Storage Element (SE) for maintaining the data storage, one or more Computing Elements (CE) for accepting and scheduling jobs and many Worker Nodes (WN) for jobs execution. In addition the clusters contain the monitoring servers, as well as User Interface (UI) allowing for interactive work of physicists.

Site:	Physical CPU	Cores	Disk capacity (TB)
FMFI	60	240	160
IEP-SAS	53	222	176

Table 1. Current parameters of Slovak clusters

For the practical reasons both clusters use the Scientific Linux OS [9] (in version 5.4 and 5.5) The middleware used at the moment is gLite [10], versions 3.1 and 3.2.

There are different technologies used for SE in Slovak clusters. The FMFI cluster is using the Disk Pool Manager (DPM) [11], whereas the IEP-SAS cluster is using the dCache [12]. In addition, both clusters have deployed the xrootd [13] storage, which is currently used by ALICE analysis jobs.

The sharing of resources is equal for both VO (ATLAS and ALICE), implemented statically at FMFI cluster (fixed dedicated CE and WN for every VO) or dynamically (dedicated CE and shared WN) at IEP-SAS cluster.

The accounting for the last year (2010) is quite encouraging. For instance the ALICE VO has successfully processed nearly 35k jobs at FMFI cluster (up to end of October). At IEP-SAS 49k jobs was completed during two months (September and October), which was possible thanks to dynamic sharing. For various technical reasons, the priority was temporary assigned to ALICE tasks in autumn 2010, so the full cluster capacity was used for ALICE VO. Peak computing capacity reaches up to 120 jobs per day.

In order to formalize the status of Slovak clusters within WLCG project, the signing of Memorandum of Understanding is planned. This way, the clusters should also became a part of Slovak National Grid Initiative.

4 Acknowledgments

The article is dedicated to the memory of Pavel Stavina, whose enthusiasm and professional skills made many of the described things possible.

We would like to express our thanks to government of Slovak Republic for their support of the particle physics in Slovakia and our membership at CERN, which includes also funding the Slovak WLCG clusters.

We are grateful to Department of Science and Research of Ministry of Education of the Slovak Republic for their support. The work was supported in part by VEGA grant 2/0061/08 of Slovak Academy of Sciences and Ministry of Education.

We would like to thank to Institute of Informatics, Slovak Academy of Science, Bratislava, Faculty of Mathematics, Physics and Informatics UK Bratislava for the general support of grid deployment in Slovakia.

References

1. Bruning, O., et al.: LHC Design Report, v.1, CERN-2004-003-V-1, Geneva, CERN, 2004
   http://lhc.web.cern.ch/lhc/
2. ATLAS Collaboration: The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003
   http://atlas.web.cern.ch/
3. ALICE Collaboration: The ALICE Experiment at the CERN LHC, JINST 3 (2008) S08002
   http://alice.web.cern.ch/
4. Proposal for Building the LHC Computing Environment at CERN, CERN/2379/Rev., Geneva, CERN, 2001
   http://lcg.web.cern.ch/lcg/
5. http://www.eu-egee.org/
6. http://www.opensciencegrid.org/
7. http://www.ndgf.org/ndgfweb/home.html
8. http://gstat-prod.cern.ch/gstat/site/FMPhI-UNIBA/
   http://gstat-prod.cern.ch/gstat/site/IEPSAS-Kosice/
9. http://linux.web.cern.ch/linux/
10. http://glite.web.cern.ch/glite/
11. https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm
12. http://www.dcache.org/
13. http://xrootd.slac.stanford.edu/