ATLAS-EDG Task Force
status report

Oxana Smirnova
LCG/ATLAS/Lund
oxana.smirnova@cern.ch
September 3, 2002, Budapest
5th EU DataGrid Conference

Members and sympathizers

ATLAS		EDG
Jean-Jacques Blaising	*Laura Perini*	*Ingo Augustin*
Frederic Brochu	*Gilbert Poulard*	*Stephen Burke*
Alessandro De Salvo	*Alois Putzer*	*Frank Harris*
Michael Gardner	*Di Qing*	*Bob Jones*
Luc Goossens	*David Rebatto*	*Peter Kunszt*
Marcus Hardt	*Zhongliang Ren*	*Emanuele Leonardi*
Roger Jones	*Silvia Resconi*	*Charles Loomis*
Christos Kanellopoulos	*Oxana Smirnova*	*Mario Reale*
Guido Negri	*Stan Thompson*	*Markus Schulz*
Fairouz Ohlsson-Malek	*Luca Vaccarossa*	*Jeffrey Templon*
Steve O'Neale		*and counting…*

The short-term use case

- ATLAS is eager to use Grid tools for the Data Challenges
  - ATLAS Data Challenges are already on the Grid (NorduGrid, USA)
  - The DC1/phase2 (to start in October) is expected to be done mostly using the Grid tools
- By September 16 (ATLAS SW Week), we should try to evaluate the usability of EDG for the DC tasks
- The task: to process 5 input partitions of the Dataset 2000 at the EDG Testbed + one non-EDG site (Karlsruhe)

Task description

- Input: set of generated events as ROOT files (each *input partition* ca 1.8 GB, 100.000 events); master copies are stored in CERN CASTOR
- Processing: ATLAS detector simulation using a pre-installed software release 3.2.1
  - Each input partition is processed by 20 jobs (5000 events each)
  - Full simulation is applied only to filtered events, ca 450 per job
  - A full event simulation takes ca 150 seconds per event on a 1GHz processor)
- Output: simulated events are stored in ZEBRA files (ca 1 GB each *output partition*); an HBOOK histogram file and a log-file (stdout+stderr) are also produced. 20% of output to be stored in CASTOR.
- Total: 9 GB of input, 2000 CPU-hours of processing, 100 GB of output.

Execution of jobs

- It is expected that we make full use of the Resource Broker functionality
  - Data-driven job steering
  - Best available resources otherwise
- A job consists of the standard DC1 shell-script, very much the way it is done in a non-Grid world
- A Job Definition Language is used to wrap up the job, specifying:
  - The executable file (script)
  - Input data
  - Files to be retrieved manually by the user
  - Optionally, other attributes (maxCPU, Rank etc)
- Storage and registration of output files is a part of the job

What's already there

- ATLAS 3.2.1 RPMs are distributed with the EDG tools to provide the ATLAS runtime environment
- Validation of the ATLAS runtime environment by submitting a short (100 input events) DC1 job was done at several sites:
  - CERN
  - NIKHEF
  - RAL
  - CNAF
  - Lyon
  - Karlsruhe – in progress
- A fruitful cooperation between ATLAS users and EDG experts
- The task force attributes:
  - Mailing list atlas-edg@cern.ch
  - Web-page http://cern.ch/smirnova/atlas-edg

What's almost there

- Input file replication: apparently, not a trivial procedure, requiring:
  - Launching a job to retrieve files from CASTOR to an SE (rfcp is not available on UI, GridFTP is not installed on CASTOR)
  - Executing a full bunch of GDMP directives
    - Define the GDMP_CONFIG_FILE depending on user's VO
    - Register files in the local catalogue
    - Check mutual subscription of SE's
    - Publish the catalogue
    - Initiate replication
- Theoretically, it works; practically, there's always something broken along the chain (e.g., RC can not be updated or can not be found altogether)
- So far, the recommended procedure worked for NIKHEF (input partitions 0003 and 0004)

What's not yet there

- Submission of long jobs
  - A job exceeding approx. 20 minutes is often being "lost" by the system: the returned status is wrong, result retrieval using the EDG tools is impossible
  - A temporary solution: the production testbed has a "regular" RB for frequent submission of short jobs, and an "ATLAS" RB for long jobs (sees only CERN and Karlsruhe sites, to be extended further)
  - This solution works sometimes (input partitions 0001 and 0002 are successfully processed), but it's not feasible in a long run

Other problems applications will face

- Installation of applications software
  - *should not be combined* with the system installation
- Authentication & authorization, users and services
  - *the procedure of accepting new CAs has to be streamlined* (40-odd countries to accommodate)
- Documentation
  - Has to be more user-oriented and concise
- Information system
  - despite being LDAP-based, lacks hierarchy
  - static to big extent
  - very difficult to browse/search and retrieve relevant info
- Data management
  - information about existing file collections is not easy to find
  - management of output data is mostly manual (can not be done via JDL)

Summary

- ATLAS users are ready and waiting
  - The Testbed is basically ready
  - 4 volunteers are named
- As soon as one real job runs and is validated, the rest 99 are expected to follow smoothly
  - Given the production Testbed's CPU power, it will take just one day
- Hopefully, this will happen before September 19
- Next meeting of the Task Force: 9/11, 10:00 at CERN (40-R-D10)