

# Computing report, ATLAS/Geneva

CHIPP - LHC computing and analysis workshop  
CSCS, Manno, 25/8-26/8 2005

Frederik Orellana  
University of Geneva/CERN

- Where are we in Geneva?
  - Middleware
  - Resources
  - Operation
- Outlook
- Issues
- Use case: Rome data reconstruction
- Conclusion

- General considerations
  - Harness as many available CPU cycles as possible
  - Do it with as few human CPU cycles as possible
  - Work towards a homogeneous Swiss ATLAS infrastructure
  - Be ready for overloaded CERN infrastructure and unavailable data
  - Implement simple user administration and prioritisation system

- NorduGrid
  - Deployable on small sites like ours (*one* server)
  - Easy installation/configuration
  - Available on a variety of Linux distros
  - Non-intrusive
  - → Geneva site part of country-wide infrastructure

Grid Monitor - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://www.nordugrid.org/monitor/

Search Misc Info TechNews SciNews News Ferme de calcul "M... x3 FO Google UF

### Grid Monitor

Norway	IBM 1300 cluster - FI>	44	0+0	0+0
	Oslo Grid Cluster	33	0+2	0+0
	UIO Grid	57	0+56	0+111
Russia	Alice SPbSU Cluster	10	2+0	0+0
	The SPbSU First Clust>	1	0+0	0+0
Slovenia	SIGNET	40	0+0	0+0
SriLanka	SweLanka LK	10	0+0	0+0
Sweden	Bluesmoke (Swegrid, NS>	95	94+0	13+0
	Cymbal	3	0+0	0+0
	Green	74	0+67	28+6
	Hagrid (SweGrid, Uppm>	100	82+0	26+0
	Hive (Swegrid, UNICC)	100	90+0	0+0
	Ingrid (SweGrid, HPC2N)	93	0+60 (no queue info)	26+0
	Kosufy farm	50	0+2	761+0
	Monolith (NSC)	388	0+328 (queue down)	1+146
	Quark Cluster	7	0+0	0+0
	Sigrd (SweGrid, Luna>	99	37+62	17+14
Switzerland	SweLanka SE	10	0+0	0+0
	Bern ATLAS Cluster	12	0+0	0+0
	PHOENIX (CSCS)	28	0+28	7+0
	Test cluster at DPNC	2	0+0	0+0
	UBELIX (University of>	96	0+82	0+9
<b>TOTAL</b>	<b>51 sites</b>	<b>4784</b>	<b>312 + 3323</b>	<b>4377 + 934</b>

Done

## Middleware



- User administration
  - Country-wide directory of high energy physicists
  - Each institute maintains the user records via a web form
  - Making a user a member of the Swiss VO is simply done via this form

CHIPP Members Directory - Mozilla Firefox <2>

http://www.chipp.ch/directory/inde

Search Misc Info TechNews SciNews News Ferme de calcul "M..." x3 FO

CHIPP Members Directory

Address book - view

<b>Full Name:</b>	Christian Häberli
<b>First name:</b>	Christian
<b>Last name:</b>	Häberli
<b>Public Key:</b>	/C=CH/O=CERN/OU=GRID/CN=Christian Haerberli 5730
<b>Institution:</b>	Universität Bern
<b>Group:</b>	ATLAS
<b>Position:</b>	Postdoc
<b>Phone:</b>	0316314062

Done

All Institutions
Universität Zürich - Theorie
Université de Fribourg
All members
P
EB
 One page

First name	Last name	Institution	Group	Position	Phone	EMail	Board	View
Christian	Häberli	Universität Bern	ATLAS	Postdoc	0316314062	Christian.Haerberli@cern.ch	CB	<a href="#">View</a>

http://www.chipp.ch/directory/index.php?menuaction=addressbook.chaddressbook.view&ab\_id=128&sessionid=43401bfa6c5225f57c935f7448...

## Middleware



- Batch system: PBS
  - Configured to give priority to Swiss VO members
  - Easily deployable with client setup script



- Old test cluster at the DPNC
  - Server: grid01.unige.ch, 300 GB storage via GRIDFTP
  - 3 attached worker nodes: scpc24, scpc27, scpc30.  
Desktop machines, grid jobs run at nice 19
- New cluster at the DINF
  - Server: grid00.unige.ch, 9.6 TB storage
  - 12 worker nodes, SUN Fire V20z, 2 x AMD Opteron, 2.4 MHz, 4 GB RAM

## Resources



- New hardware has been tested with Scientific Linux, ATLAS SW and NorduGrid, all ran fine
- The system is expected ready in ~two weeks
- Presumably the system REMBO (UniGe/DINF) will be used for installation and maintenance
- Questions to sysadmins:
  - Cluster management systems?
  - Disk layout/partitioning?
  - Sharing of directories with experiment's software, NG session directories: NFS, NCP, SMB?

## Operation



- Currently one 10% system administrator for the old test cluster
- Useful exercise: perfectly possible to run a grid site without dedicated manpower
- Main problem: update, debug and patch ATLAS software


- More users
  - Set up more virtual organisations
- More manpower
  - FO: 10% → 100%
  - DINF cluster is hosted in computing centre
- More resources
  - Extend DINF cluster → 75 nodes
  - Deploy PBS clients on more Linux desktops wherever possible

## New functionality GUI for submission and monitoring

The screenshot shows the 'AtCom - an Atlas Commander' window. It features a table with columns for Job Name, Job ID, Job Status, CS, Host, AMI Status, and User. The table lists various jobs with their current statuses, such as 'Submitted', 'Validated', 'Undecided', 'Not\_Fou...', and 'Done'. On the right side, there are radio buttons for 'View all jobs', 'View only running jobs', and 'View only done jobs'. Below these are buttons for 'Load jobs from AMI' and 'Clear'. A 'Jobs statistics' section contains a vertical bar chart with three segments: 'Wait' (black), 'Run' (blue), and 'Done' (green). At the bottom, there are buttons for 'Monitor' and 'Logs', and a status bar showing 'Validation of RecExTB\_Combined\_912\_2102286.00006 done : Validated (0 jobs in the queue)'.

Job Name	Job ID	Job Status	CS	Host	AMI Status	User
RecExTB_Combined_912_2102245.00...	638352	RUN	LSF@CERN	lxb1053	Submitted	ctbprod
RecExTB_Combined_912_2102245.00...	638354	RUN	LSF@CERN	lxb0416	Submitted	ctbprod
RecExTB_Combined_912_2102245.00...	638355	RUN	LSF@CERN	lxb1078	Submitted	ctbprod
RecExTB_Combined_912_2102245.00...	638357	RUN	LSF@CERN	lxb1666	Submitted	ctbprod
RecExTB_Combined_912_2102245.00...	638358	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102245.00...	638360	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102246.00...	618298	NOT_FOU...	LSF@CERN		Validated	ctbprod
RecExTB_Combined_912_2102246.00...	638361	DONE	LSF@CERN		Validated	ctbprod
RecExTB_Combined_912_2102246.00...	638364	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102249.00...	618299	DONE	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102249.00...	638366	RUN	LSF@CERN	lxb0276	Submitted	ctbprod
RecExTB_Combined_912_2102249.00...	638367	RUN	LSF@CERN	lxb0490	Submitted	ctbprod
RecExTB_Combined_912_2102249.00...	638369	RUN	LSF@CERN	lxb0490	Submitted	ctbprod
RecExTB_Combined_912_2102249.00...	638372	RUN	LSF@CERN	lxb1094	Submitted	ctbprod
RecExTB_Combined_912_2102249.00...	638374	RUN	LSF@CERN	lxb1076	Submitted	ctbprod
RecExTB_Combined_912_2102249.00...	638376	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102249.00...	638378	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102261.00...	638380	RUN	LSF@CERN	lxb0684	Submitted	ctbprod
RecExTB_Combined_912_2102273.00...	618303	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102273.00...	638383	RUN	LSF@CERN	lxb0438	Submitted	ctbprod
RecExTB_Combined_912_2102273.00...	638385	RUN	LSF@CERN	lxb0482	Submitted	ctbprod
RecExTB_Combined_912_2102273.00...	638388	RUN	LSF@CERN	lxb1824	Submitted	ctbprod
RecExTB_Combined_912_2102273.00...	638389	RUN	LSF@CERN	lxb0684	Submitted	ctbprod
RecExTB_Combined_912_2102273.00...	638391	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102283.00...	638392	NOT_FOU...	LSF@CERN		Undecided	ctbprod
RecExTB_Combined_912_2102286.00...	618305	NOT_FOU...	LSF@CERN		Validated	ctbprod
RecExTB_Combined_912_2102287.00...	618307	NOT_FOU...	LSF@CERN		Validated	ctbprod
RecExTB_Combined_912_2102287.00...	618308	NOT_FOU...	LSF@CERN		Validated	ctbprod
RecExTB_Combined_912_2102287.00...	638394	RUN	LSF@CERN	lxb0061	Submitted	ctbprod
RecExTB_Combined_912_2102287.00...	638396	RUN	LSF@CERN	lxb0401	Submitted	ctbprod
RecExTB_Combined_912_2102287.00...	638406		LSF@CERN		Submitted	ctbprod
RecExTB_Combined_912_2102287.00...	638415		LSF@CERN		Submitted	ctbprod
RecExTB_Combined_912_2102292.00...	638559		LSF@CERN		Submitted	ctbprod
RecExTB_Combined_912_2102292.00...	638565		LSF@CERN		Submitted	ctbprod
RecExTB_Combined_912_2102294.00...	638570		LSF@CERN		Submitted	ctbprod
RecExTB_Combined_912_2102289.00...	618309		LSF@CERN		Submitted	ctbprod

## Issues




- Preparing and submitting jobs still needs the assistance of an 'expert'
  - Main problems
    - difficulties in getting the data
    - ATHENA
- ATLAS: meta-data catalogue
- ATLAS: reliable file transfer service
- Tape/backup strategy
- Shared file system: NFS may not scale

## Use case: Rome data reconstruction



- Job preparation and submission
  - Running on 2000 files necessitated some automation: script with submission loop and job babysitting
  - NorduGrid XRSL template with placeholders for input/output files, etc. parsed by script, resulting in 2000 XRSL files
  - File catalogue (RLS) could have been used, but we saw now reason

# Use case: Rome data reconstruction



## XRSL:

```
&(executable="ge-nordwrapper.sh")
(arguments="10.0.1" "_inputFile_" "Trigger_topOptions_rome_grid.py")
(stdout=atlasreco.out)
(stderr=atlasreco.err)
(inputFiles=(Trigger_topOptions_rome_grid.py "" )(jobOfragment_TrigElectronHypo.py "" )(jobOfragment_TrigSteerMonitor.py
"" )(HLTsequence.xml "" )(HLTsignature.xml "" )(_inputFile_ "_castorDir__inputFile_")(InstallArea.tar.gz "" ))
(outputFiles=( _outputFile_ "" ))
(gridtime="30 minutes")
(jobname="ATLAS dc2 Zee reconstruction, _inputFile_")
(gmlog=log)
(runTimeEnvironment=APPS/HEP/ATLAS-10.0.1)
(memory=1512)
(disk=1200)
```




## Use case: Rome data reconstruction




- Babysitting script ran well, submitting and retrieving 100 jobs.
- Ready to start production as soon as data issues have been solved
- Should be easy to modify our scripts for other use cases, see the public CVS repository on <https://savannah.cern.ch/projects/edginstall/>

# Use case: Rome data reconstruction



## Conclusion



- Geneva is well integrated in the Swiss ATLAS computing infrastructure
- In the near future we will also start contributing significant resources to the common pool
- System administration will be aligned with general university services
- ATLAS still has some issues to be fixed before massive user analysis on the grid is realistic:  
input data, Athena outside of CERN (afs, conddb)