CIMENT - Calcul Intensif, Modélisation, Expérimentation Numérique et Technologique

Un regroupement de pôles meso-info mutualisés par une grille légère CIGRI

Bruno Bzeznik, Laurent Desbat
Universités de Grenoble, UJF

https://ciment.ujf-grenoble.fr
http://cigri.imag.fr

Journée sur l'organisation et l'activité des méso-centres de calcul en France.
13 fév. 2008, IHP, Paris.

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CIMENT : des moyens de calcul au service de la science

SCCI (2005 et 2006, SUN) 160 opteron Service de Calcul Intensif de l'Observatoire de Grenoble

MIRAGE (2004, SGI) 2*16 Itanium II:
Meso Informatique Répartie pour des Application en Géophysique et Environnement

CECIC(2004, IBM) 2*16 Power V :
Centre d'Expérimentation du Calcul Intensif en Chimie

GrappePCS (2003, HP) 208 itanium II:

Grappes de PCs , recherche en informatique

BioIMAGe:
Biologie Imagerie
PHYNUM:
Physique Numérique
(2004, SGI) 2*16 Itanium II
(2006, SGI) 72 itanium II + 28 xeon 64b

HPC Ciment platforms

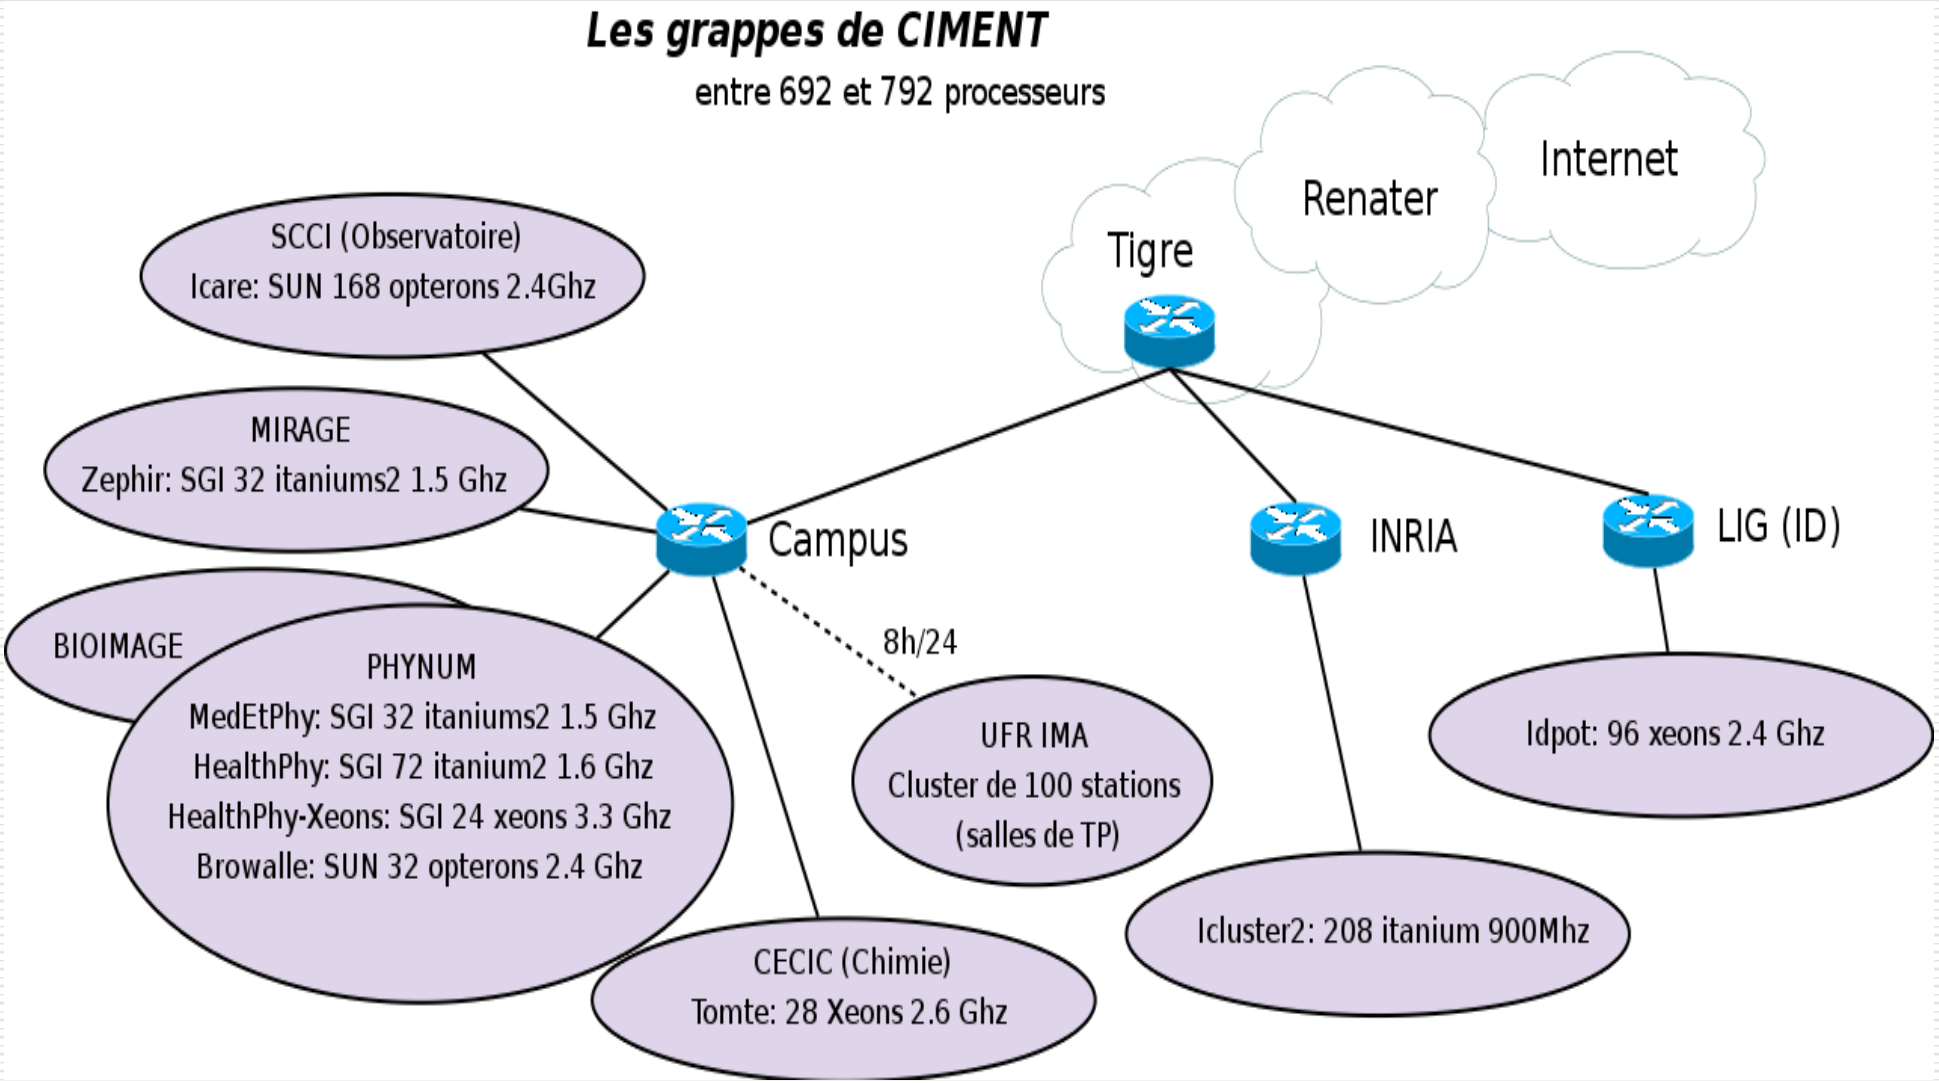

Bilan en équipement

- Une coordination d'équipement intermédiaire en calcul intensif centré sur des projets scientifiques
- Budget CPER 2000-2006

Organismes	Région RA	Métro	Ministère	INRIA	**Total**
Subventions €HT	800,357.31	571,684.00	523,918.51	381,122.54	**2,277,082.36**

- Réalisations

plateformes de calcul	**Type**	**coût en €HT**
plate forme CECIC (Chimie)	3 quadri pro IBM SP3, 1SGI0200	228673.53
Jouvence de CECIC	2 hexapro IBM P5, 1.5Ghz	163096.02
Plate forme BioIMAGe (pôle Santé)	24 bi-pro XP1800+ Athlon	105648.31
Grappe de 200PC	104 bi-pro Itanium2, HP	588061.36
Plate forme PhyNum	40 bi-pro XP2000+ Athlon	207830.04
Plate forme MIRAGE (environnement)	2 hexa-pro SGI Altix 350	227123.55
Plate forme SCCI (Observatoire)	42 quadri-pro opteron SUN couplés GRID5000	465,219.92
MedEtPhy (BioIMAGe+Phynum)	2 hexa-pro SGI Altix 350	194,437.90
HealthPhy (BioIMAGe+Phynum)	SGI Altix 450 (72 cœurs Itanium et 28 Xeon)	310600.00

Perspectives scientifiques

- CIMENT structuré sur des pôles scientifiques et leurs projets (pilotage par les chercheurs utilisateurs et les ingénieurs)
- Accompagne la restructuration de la communauté de Physique Numérique à Grenoble
  - => un axe fort vers la simulation pour les nano sciences (**minatec** et **minalogic**)
  - CIMENT labellisée **plateforme du RTRA** nanosciences aux limites de la nanoélectronique
- En liens avec les clusters
  - « informatique, signal, logiciels embarqués »
  - « matériaux »
  - ...
- Renforce les projets des laboratoires et instituts
  - Molecular universe
  - Environnement
  - Chimie, ...

Bilan de la formation

- Formation doctorale
  - 2 modules d'une trentaine d'HeTD
  - Doctorants + quelques M2
- Formation permanente et d'automne
  - Alternance CNRS et Univ. (Introduction au calcul scientifique).

Bilan d'expertise

- Expertise sur les marchés d'équipement pour le calcul (CCTP, etc.) => ex. de GRID 5000
- Expérimentation transversale : CiGri
- Partenariat ICATIS/UFRIMA/ID-IMAG pour le calcul sur PC d'enseignement couplés à CiGri
- Partenariat LIG/Bull, thèses cifre, projets communs
- Partenariat avec la DSI (ex-CICG) pour l'accueil de machines de calcul : une première expérience très réussie.

Bilan : les ingénieurs de CIMENT

- F.Berthoud (IR CNRS, PhyNum, MedEtPhy), F.Roch (IR CNRS, Obs), Pierre Vatton (IR CNRS, CECIC), Laurence Viry (IR UJF CIMENT, CRIP MIRAGE), Nicolas Capit (24 mois de CCD sur CIMENT GRID).
- Un IR UJF CIMENT 2006 : Bruno Bzeznik système/ réseau/grappes/grilles.
  - => 2 Ingénieurs CIMENT de compétences très différenciés
  - => Projets coordonnés possibles
- De nouveaux ingénieurs de pôles souhaitant développer leur activité calcul (Nathalie D'Agostino Philippe Beys, SPECTRO), etc.
- **Un réseau d'ingénieurs bien structuré prêt à une collaboration régionale sur le calcul (grille, formation, partage d'expertise, ...).**

Structure et avenir

- Comité de Pilotage (sci.+ingé)
- Conseils Scientifique (experts extérieurs + représentants des tutelles)
- Avenir : CIRA = Calcul Intensif en Rhône Alpes (regroupe la FLCHP, MUST et CIMENT) avec en particulier la construction/expérimentation de grilles régionales avec le CC de l'IN2P3 (RAGrid : Rhon'Alp'Grid)

Plan

- **The CIMENT Project**
- **The CiGri project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri: a lightweight Grid for parametric applications

Bruno.Bzeznik@imag.fr
UJF/CIMENT
LIG (Mescal Team)

CIMENT platforms: heterogeneous

- **About 700 processors with different architectures and power**
  32 bits, 64 bits, AMD, Intel, pure 64 bits (itanium2). Shared memory computers, distributed memory computers. Low latency (or not) networks, (Giga-ethernet, Infiniband, Myrinet,...) . Different operating systems (SLES, Solaris, Fedora,...)
- **Production hosts:**
  - Zephir, MedEtPhy, HealthPhy, Tomte
- **Experimental hosts** (Grid5000):
  - Idpot, Icluster2
- **Mixed hosts** (a part used in Grid5000)
  - Icare, cmserver

UFRIMA cluster: a cluster of desktop PCs

- Desktop PCs in practice rooms
- Used by students during the day
- Used as a diskless PC cluster by night
- Tools:
  - COMPUTEMODE (www.computemode.org) for the images, DHCP, PXE, NFS and the calendar
  - OAR as the batch scheduler

CiGri: **CI**MENT **GRI**d

- *Opensource* project started in 2002
- Financed by French research ministry : ACI GRID «CIGRID» (an engineer during 2 years: Nicolas Capit, working for the **I**nformatique **D**istribuée laboratory)
- Partners:
  - ID laboratory (today LIG/Mescal): development
  - Other CIMENT departments: applications
- In 2004, 100 000 jobs executed
- The OAR batch scheduler, used by Grid5000 and Ciment, was initialy developped during CiGri genesis
http://oar.imag.fr
- Today, the CIMENT grid executed near 3,5 millions of jobs

CiGri Goals

- Optimizing the use of the CIMENT clusters (which are heterogeneous)
- The **lightweight grid** concept:
  - Minimizing the grid computing problem : only **multiparametric applications**
  - Sites are using common administration rules
  - No big authentication problems (ssh and sudo are good enough for us)
  - Lighter than heavy grids like GLOBUS
- Campains may be composed of a very high number of small jobs
- Cigri was also used to experiment distributed scheduling in a computer science research context
- Localy idle resources are used by the grid (best-effort concept)

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri efficiency (1)

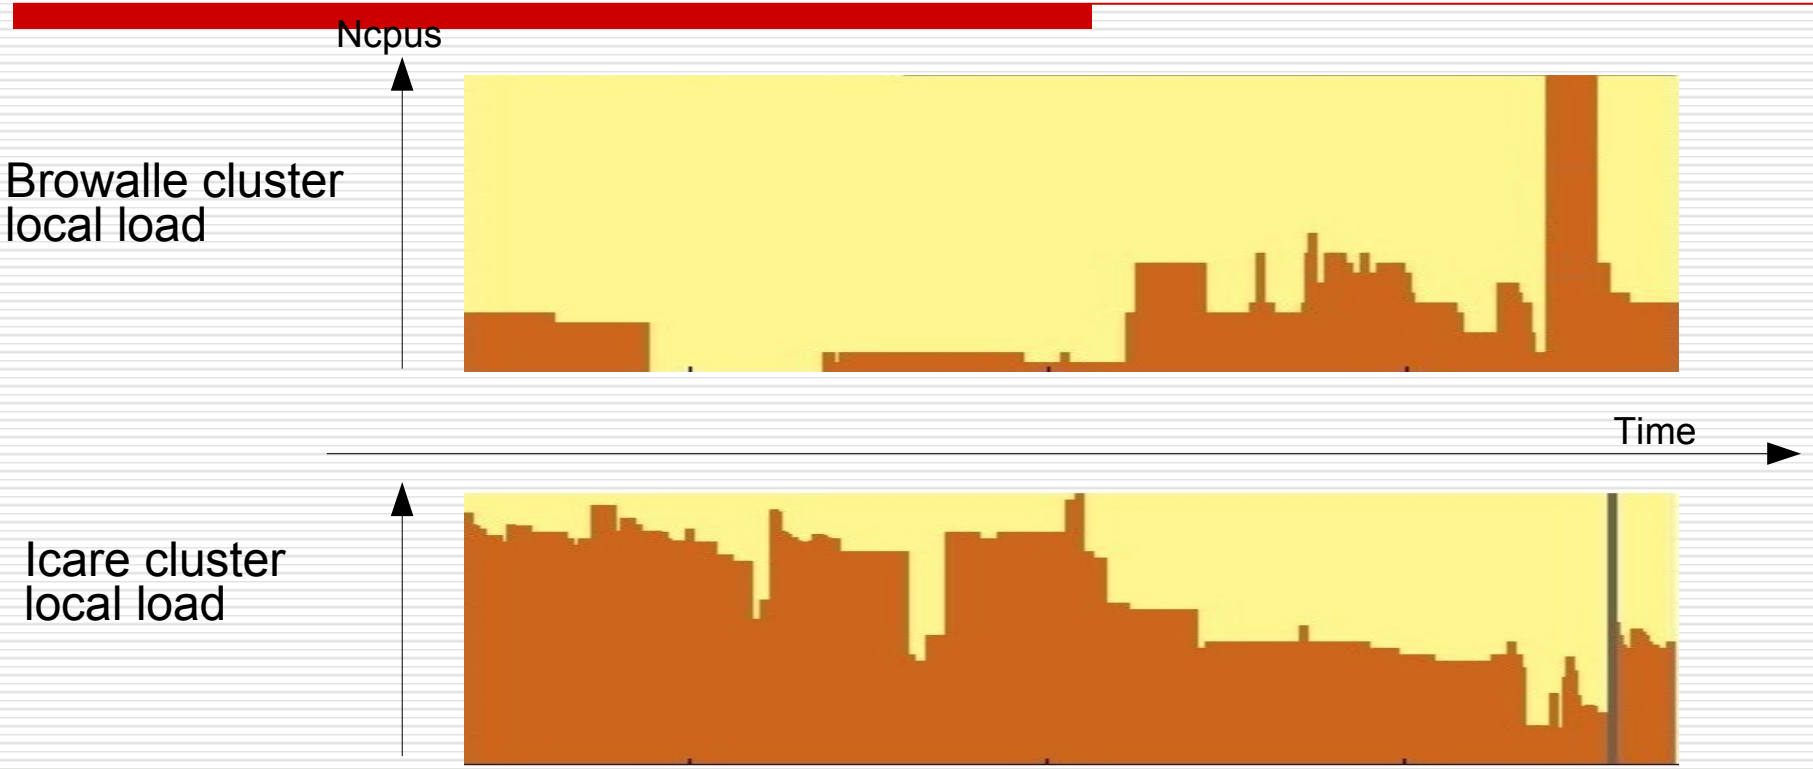

The load of the clusters is not constant and peaks are often not at the same time...

CiGri efficiency (2)

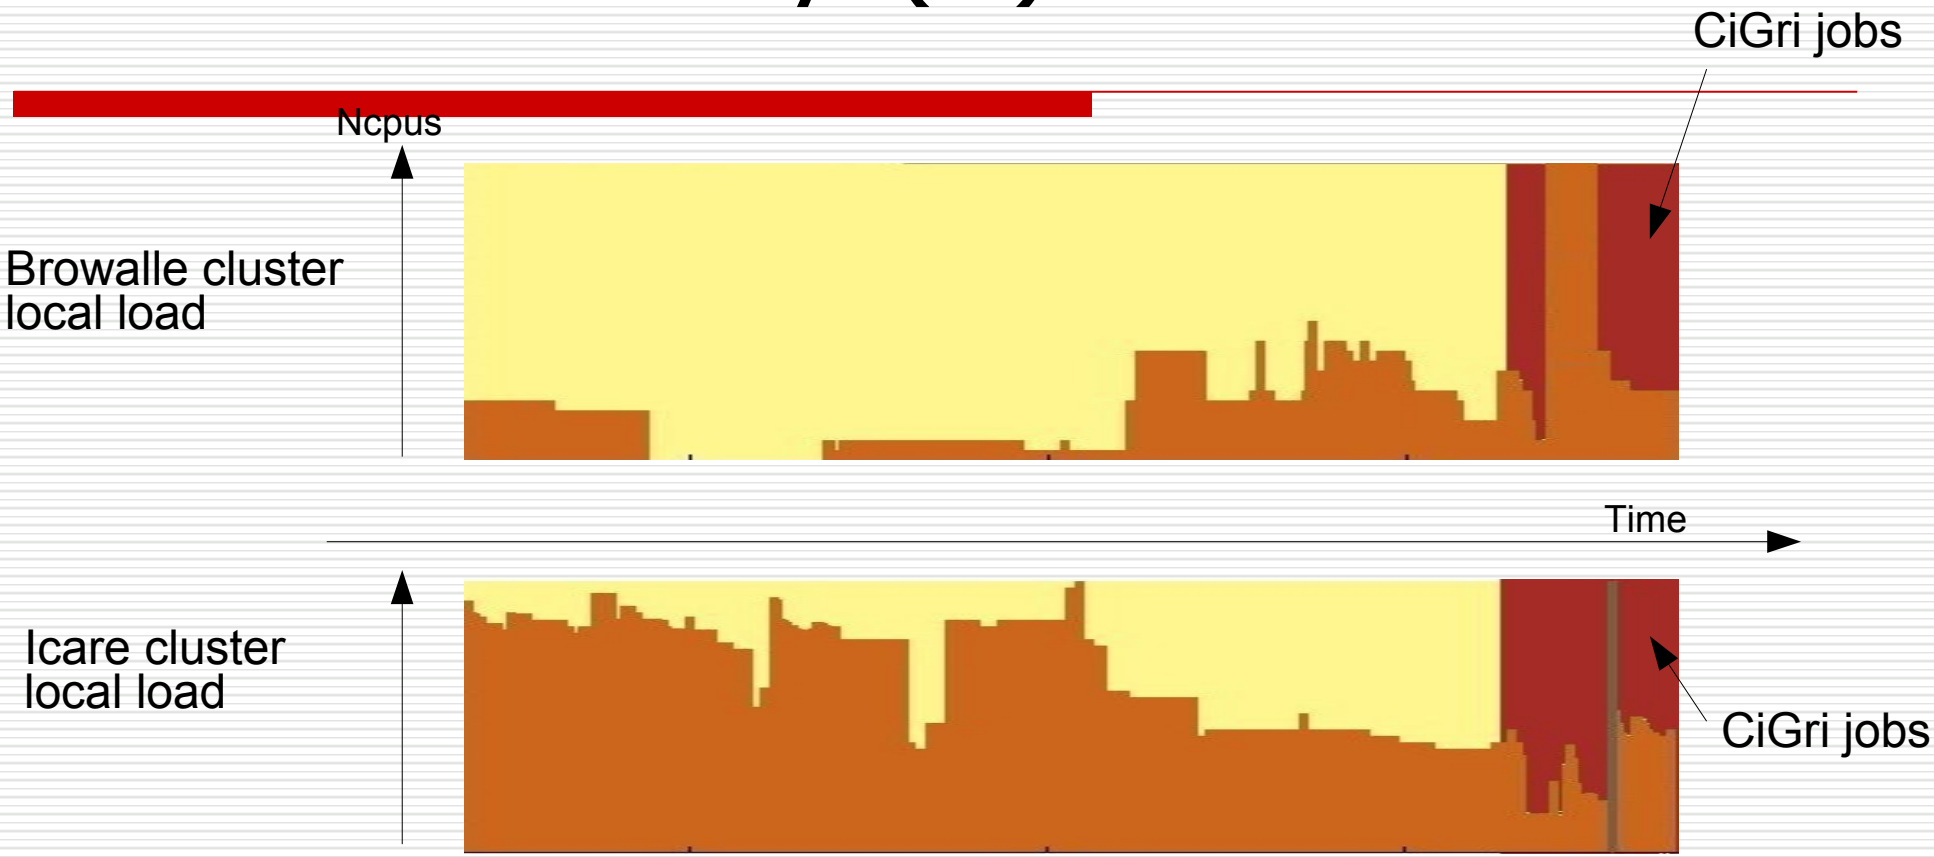

CiGri uses the idle cpus

CiGri efficiency (3)

- By tracking the users that run multiparametric applications into a set of cluster, you may:
  - localy free a cluster that needs power for other types of applications
  - may reach a 100% usage of all the clusters with a best-effort load spread
- But you need manpower to:
  - find those users
  - help them to port the application to the grid

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri architecture

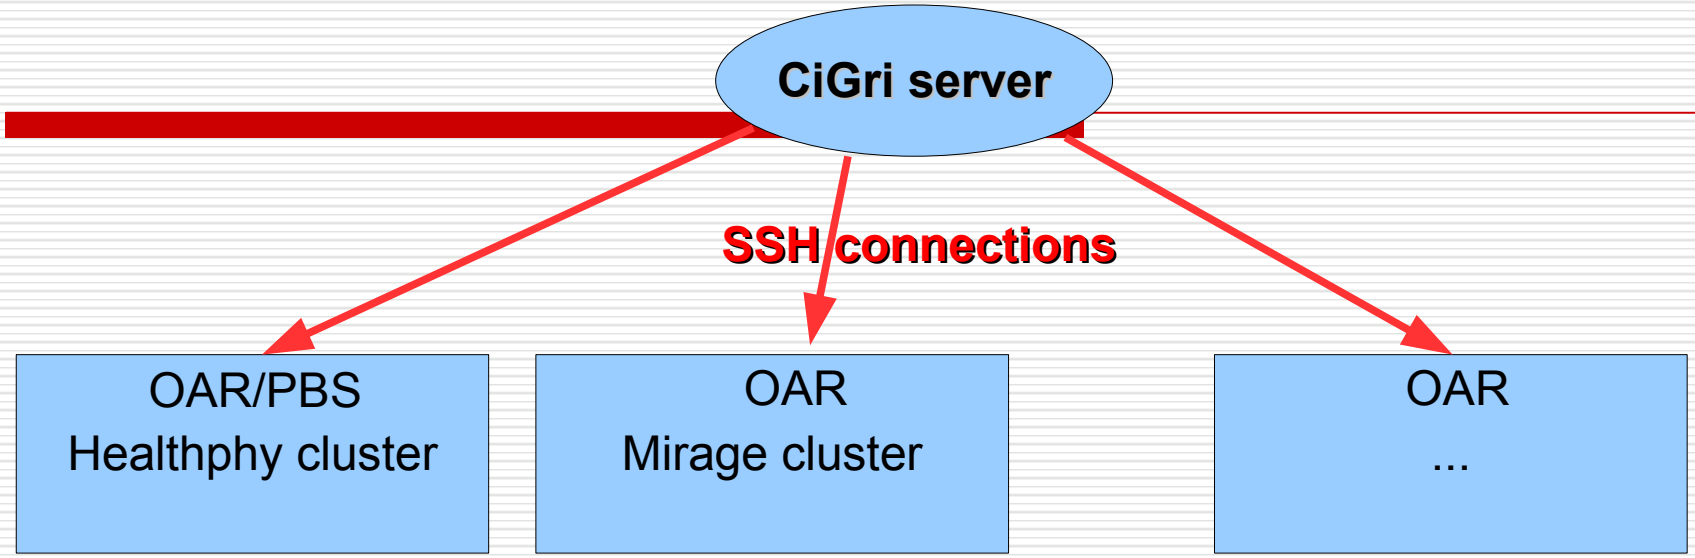

- A central CIGRI host
- Uses an SQL database as the core model
- Communicates with clusters via ssh
- **Non intrusive for local production sites**
- Submits jobs into the **OAR batch scheduler** (maybe coupled with another bs)
- Uses the « **best effort** » concept of the OAR batch scheduler (next slide)

Best-effort jobs

- Best effort jobs allow us to **exploit idle resources** of production clusters
- When a node is free, a best-effort job may be placed on it.
- If a “normal” job needs the node, the best-effort job is **killed**
- -> as a consequence, a best-effort job must be “short enough” (< 30 minutes) or there's too much chance for it to be killed
- Checkpointing may help

CiGri authentication

- CiGri accounts stored in the database and a binding with ssh accounts for each cluster of the grid

OR

- **Ldap** authentication support

And **sudo** on each cluster

Cluster requirements

- SSH access from the cigri host
- A “cigri” user
- Sudo configuration for the “cigri” user to be able to become a user of the grid
- **OAR batch scheduler (http://oar.imag.fr)**
  - It may be exclusive,
  - but may coexist with another batch system in a best-effort way (resources are “absent” when they are used by the other batch system)
- Facultative kernel level checkpoint support (BLCR, SGIGrid checkpoint,...)

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri: Features (1)

- Manages multiparametric tasks
- No communication between tasks (or only at node level)
- The user may launch thousands and thousands of tasks (jobs) in only one submission called the “campaign”
- Task's precarity is managed: if a task is killed, it's automaticaly resubmited (perhaps on another cluster)
- Error management: <>0 return code, network problems,...

CiGri: Features (2)

- Collect system for the result files (useful as a lot of files may be created)
- Jobs and campaigns may be entirely deleted
- Web interface:
  - Graphical stats
  - Submissions checking
  - Decisions taking (resubmission, stop, drop a parameter, ...)
  - Grid status
  - Forecast
- Data may be transferred with the submission

CiGri: Features that are currently in development

- Scheduling enhancement: jobs interlacing, fairsharing
- Checkpointing
- Events managing enhancements
- Perspectives: parallel jobs support, new web interface, campain initialization, deployment, ...

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri usage

- If LDAP is not used, the gridmaster creates an account on the CiGri host and asks for an account on each cluster of the grid, then creates a **binding between the grid account and the clusters accounts**.
- The user **installs and tests it's application on each cluster** of the grid that he wishes to use
- On the cigri host, he defines the campain in a **JDL file** (exemple on the next slide)
- Submission: **gridsubmit -f <JDL file>**
- Visualization into the **CiGri's web interface**
- **Collect** of the results on the cigri host

CiGri JDL file

```
DEFAULT{
    name = campagne1 ;
    # njobs = 1000;
    paramFile = param.tmp;
}
healthphy.ujf-grenoble.fr{
    execFile = /home/ciment/bzizou/test_ia64.sh ;
    execDir = /home/ciment/bzizou ;
    walltime = 00:15:00 ;
}
cmserver.imag.fr{
    execFile = /users/nis/bzeznik/test.sh ;
    checkpoint_type = blcr ;
}
```

CiGri web interface

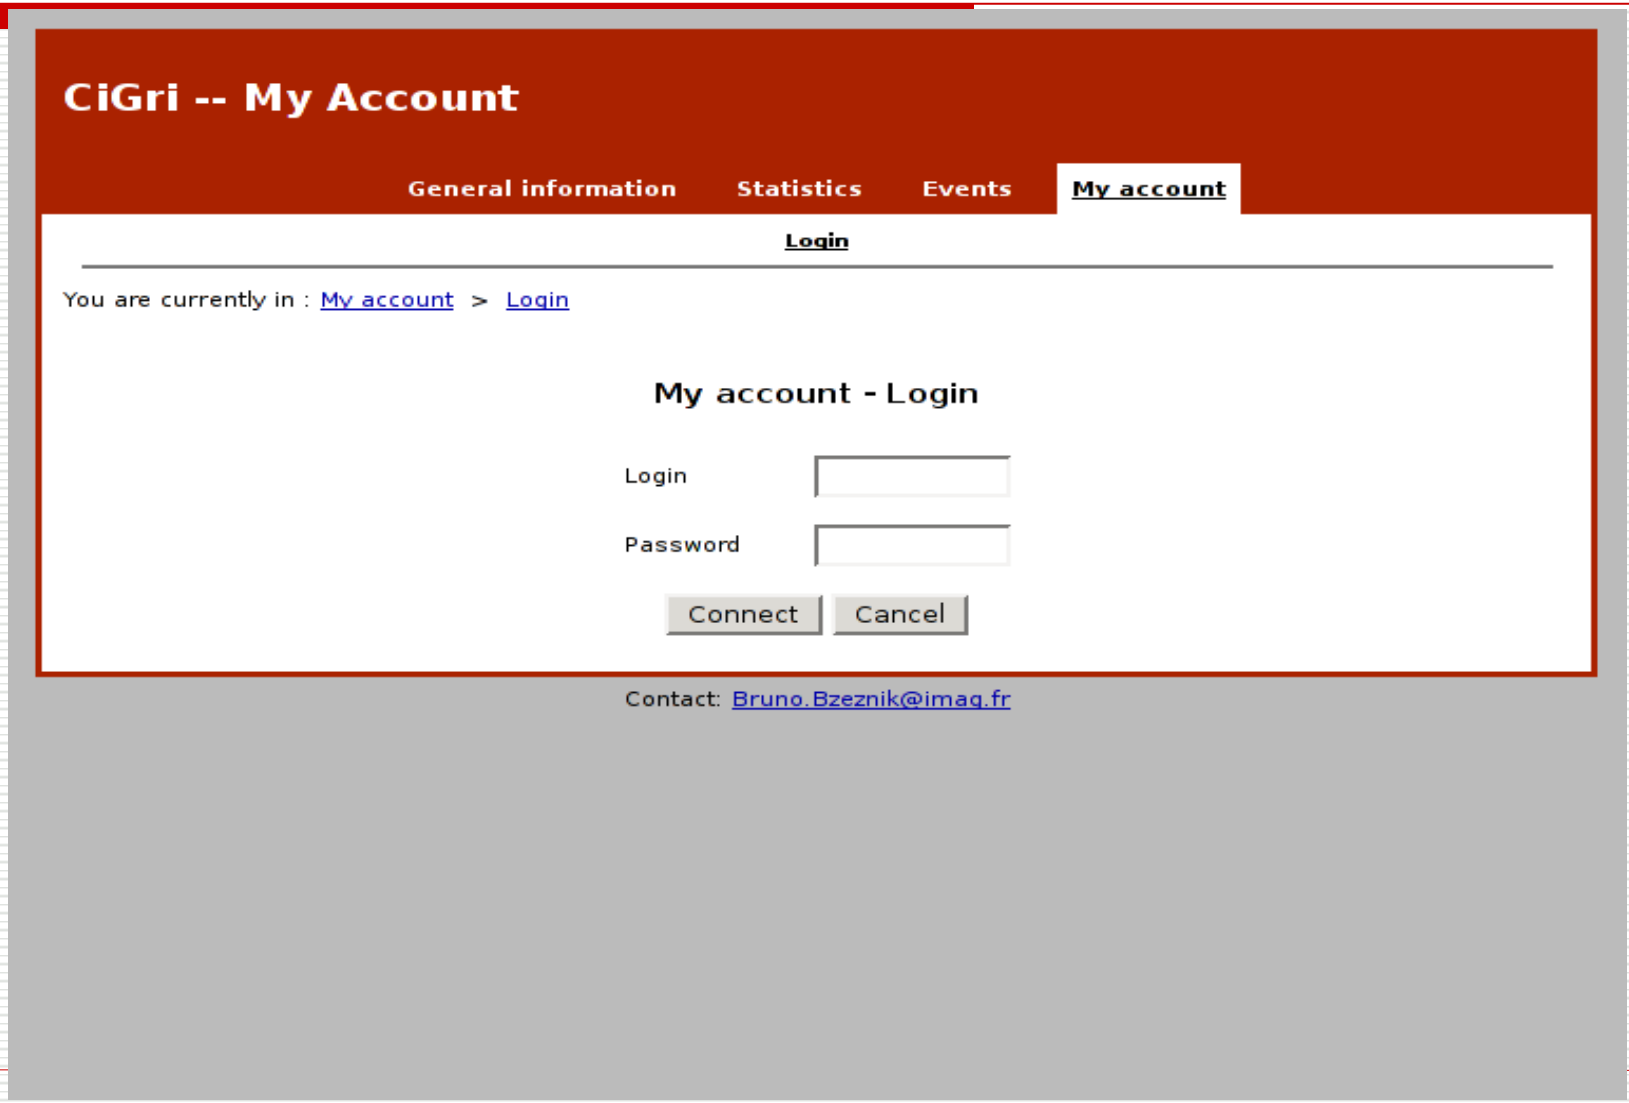

Current grid status

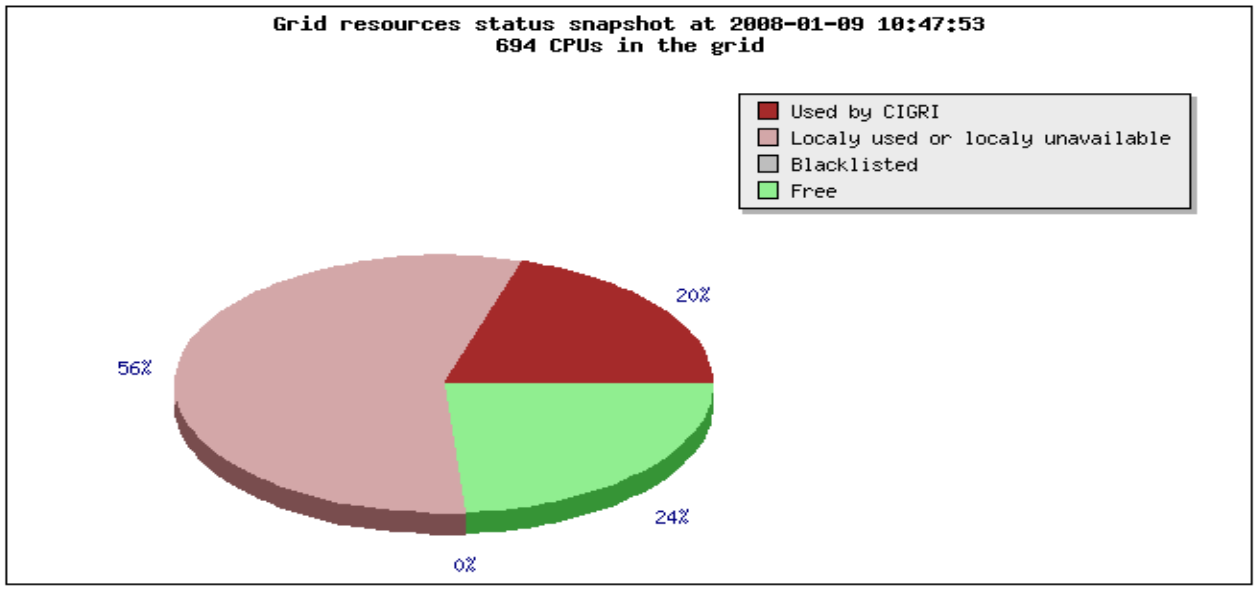

Latest status recorded on 2008-01-09 10:47:53 :

Cluster	Blacklisted	Max resources	Used resources (by cigri)	Localy used or unavailable resources	Free resources
icare.obs.ujf-grenoble.fr	no	168	91	77	0
cmserver.e-ima.ujf-grenoble.fr	no	126	2	124	0
healthphy.ujf-grenoble.fr	no	100	3	97	0
icluster2.imag.fr	no	100	0	15	85
idpot.imag.fr	no	48	0	2	46
p5cecic.ujf-grenoble.fr	no	32	0	32	0
zephir.mirage.ujf-grenoble.fr	no	32	0	28	4
medetphy.imag.fr	no	30	0	0	30
browalle.ujf-grenoble.fr	no	30	23	7	0
tomte.ujf-grenoble.fr	no	28	22	6	0
TOTAL 10 clusters	0 unavailable cluster(s)	694	141	388	165

Current multiple-jobs :


CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Status | Errors | admin: logout
You are currently in : My account > MultiJobs > MultiJob #377 > Running jobs
MultiJob #377 - Running
FORECAST: Avg: 132 / Stddev: 3.74 / Troughput: 61.71 j/h / End: in 0d,0h,6m,38s (2008-01-09 10:58:05)
STATUS: Term: 3 / Run: 141 / RemoteWait: 0 / Wait: 397 / resubmissions: 0%
Running Jobs - Executed Jobs - Waiting Parameters
Running Jobs 1 - 20 out of 141
<< < Page 1 / 8 with 20 items per page Update > >>
Job # Job name Submission date Cluster
673553 landscape_0143.png 2008-01-09 10:50:47 healthphy.ujf-grenoble.fr
673552 landscape_0142.png 2008-01-09 10:50:46 healthphy.ujf-grenoble.fr
673551 landscape_0141.png 2008-01-09 10:50:44 healthphy.ujf-grenoble.fr
673550 landscape_0140.png 2008-01-09 10:49:30 icare.obs.ujf-grenoble.fr
673549 landscape_0139.png 2008-01-09 10:49:29 icare.obs.ujf-grenoble.fr
673548 landscape_0138.png 2008-01-09 10:49:28 icare.obs.ujf-grenoble.fr
673547 landscape_0137.png 2008-01-09 10:49:27 icare.obs.ujf-grenoble.fr
673546 landscape_0136.png 2008-01-09 10:49:26 icare.obs.ujf-grenoble.fr
673545 landscape_0135.png 2008-01-09 10:49:25 icare.obs.ujf-grenoble.fr
673544 landscape_0134.png 2008-01-09 10:49:23 icare.obs.ujf-grenoble.fr
673543 landscape_0133.png 2008-01-09 10:49:22 icare.obs.ujf-grenoble.fr
673542 landscape_0132.png 2008-01-09 10:49:21 icare.obs.ujf-grenoble.fr
673541 landscape_0131.png 2008-01-09 10:49:20 icare.obs.ujf-grenoble.fr
673540 landscape_0130.png 2008-01-09 10:49:19 icare.obs.ujf-grenoble.fr

Plan

- **The CIMENT Project**
- **CiGri goals**
- **CiGri efficiency**
- **CiGri architecture**
- **CiGri features**
- **CiGri usage**
- **User examples**

CiGri: User examples

GATE (GEANT4 Application for Tomographic Emission)

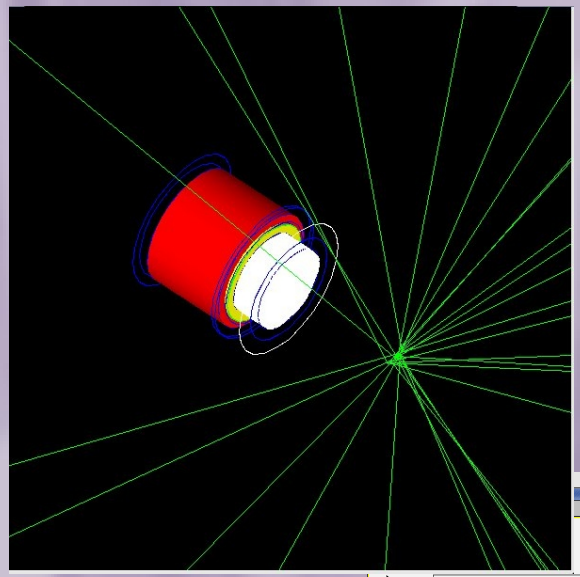

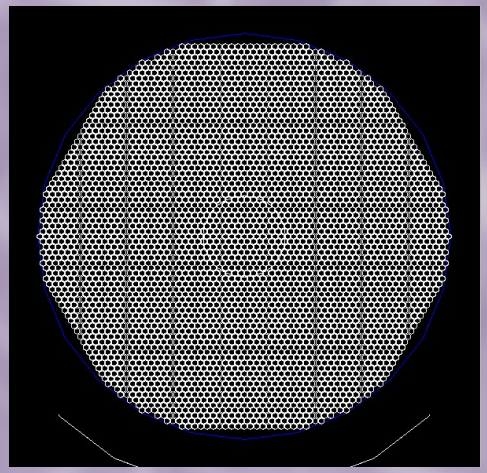

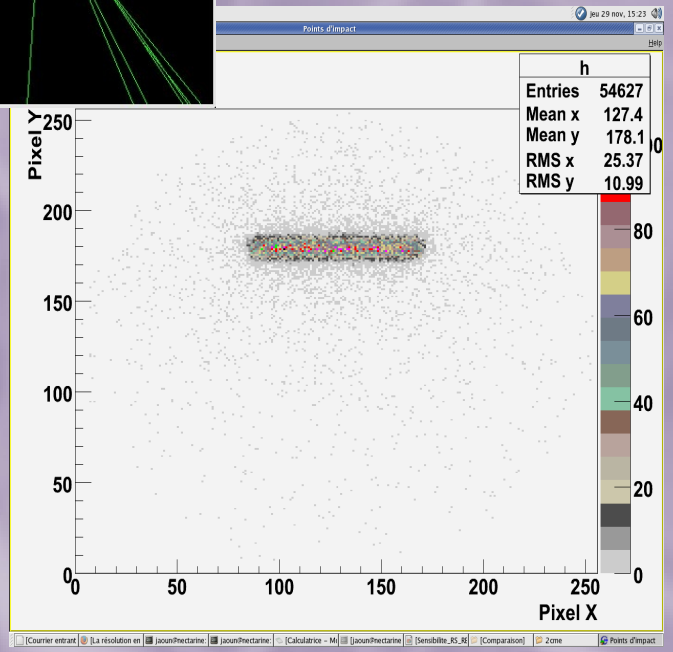

Détection de photons (1 photon détecté pour 10000 générés)
-> concaténation du résultat de nombreuses simulations

Heures de calcul: travail en cours.
2,4 années de temps CPU en 2007

CiGri: User examples

Ajustement simultané d'observations possibles de disques protoplanétaires.
Utilisation d'algorithmes génétiques pour l'exploration d'un espace de paramètres.

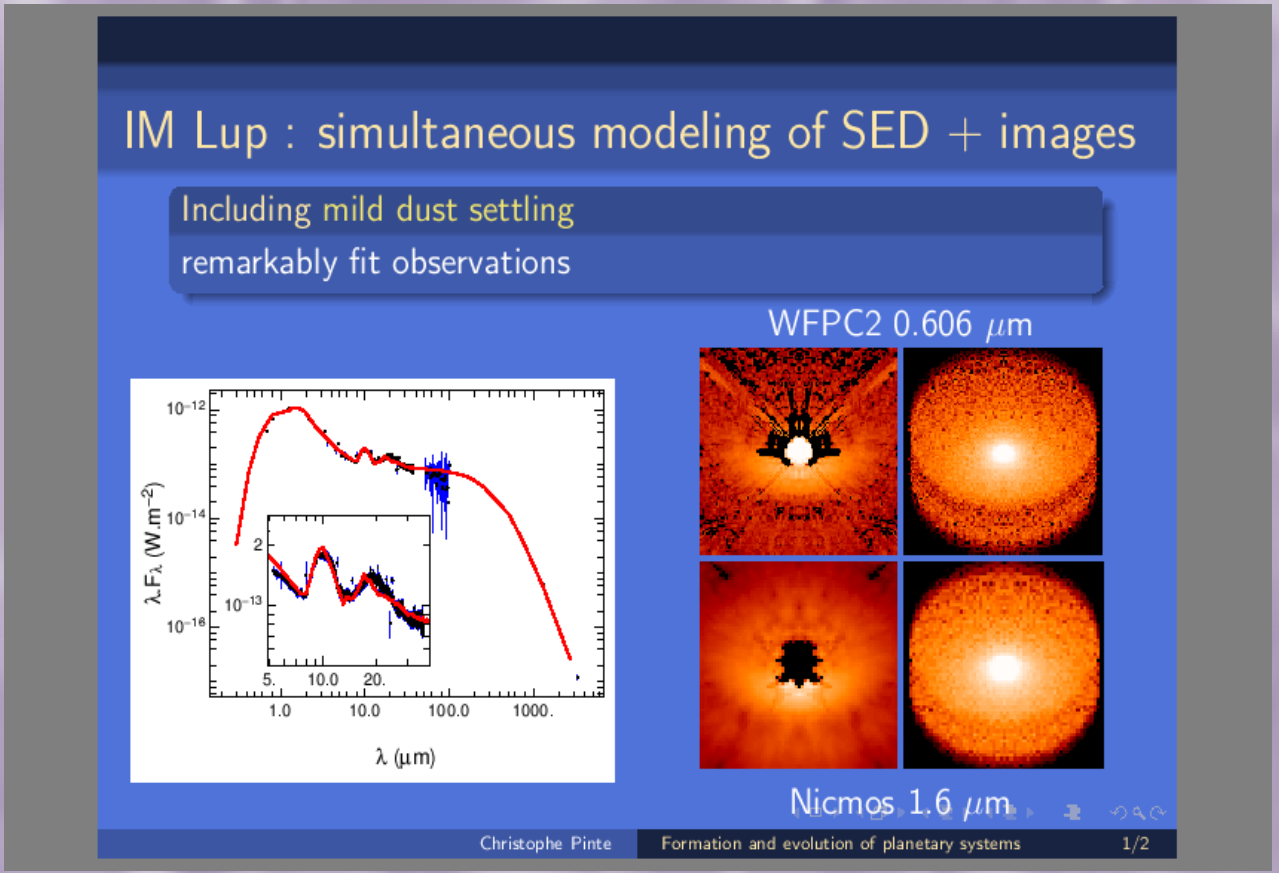

8,6 années de temps CPU utilisées en 2007

CiGri: User examples

Calculer la production ionique par les rayons cosmiques dans l'atmosphère de Titan.

Lancement d'un programme (planetocosmic - Monte-Carlo) pour un grand nombre d'énergies en parallèle, et ainsi une convolution avec un spectre de précipitations de rayons cosmiques nous donnera la production.

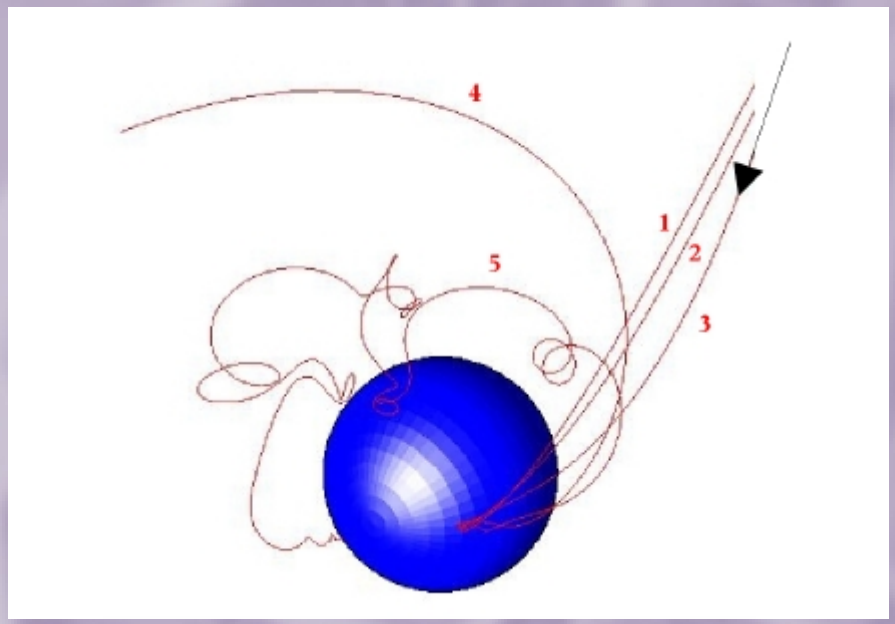

CiGri: User examples

Modélisation hydro-météorologique en Afrique de l'ouest (code MAR)
Couplage sol-végétation-atmosphère

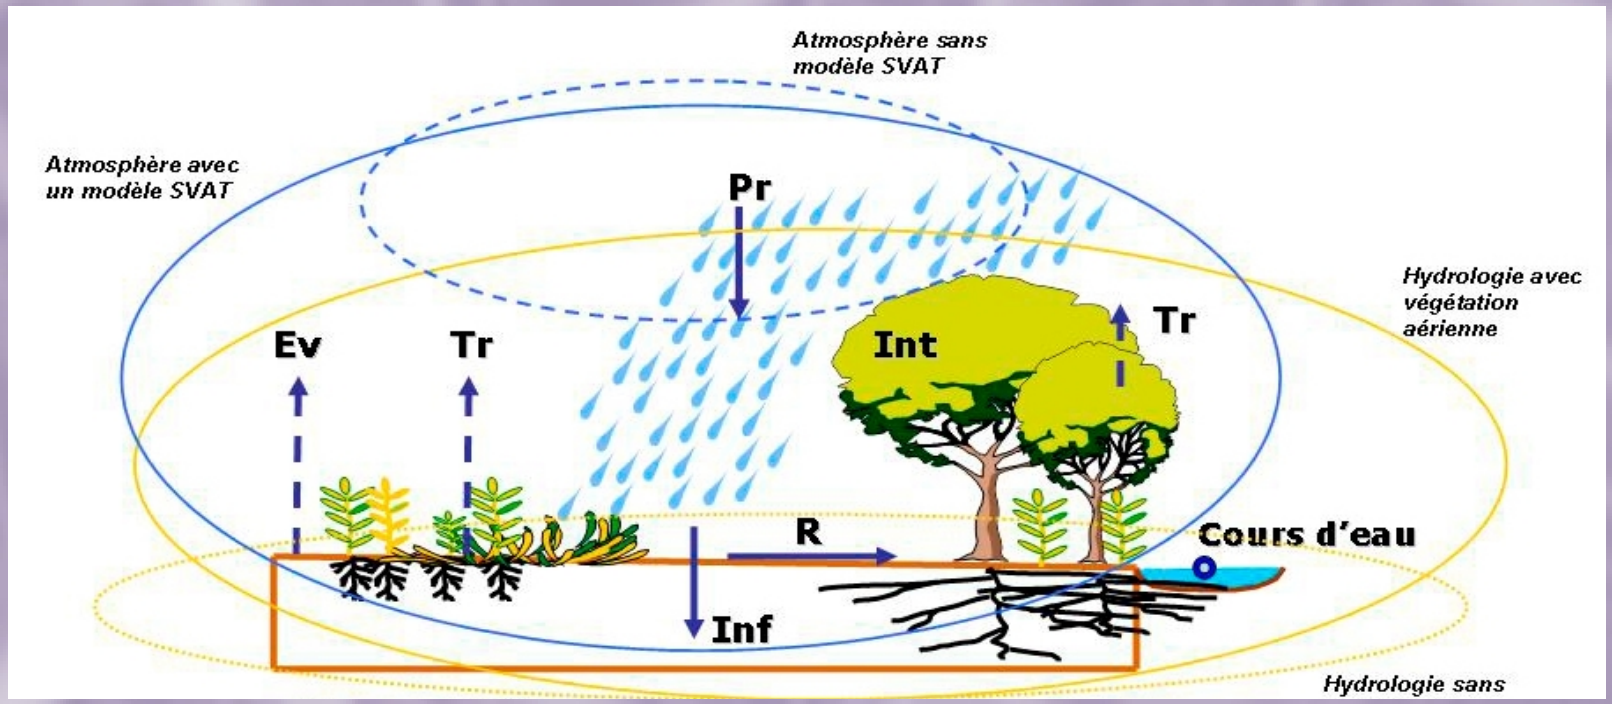

Mise au point des campagnes en cours. Nécéssité de checkpointing dans CiGri (jobs de plusieurs heures). Experimentations prévues avec DIET.

CiGri: User examples

Calculs de trajectoires moléculaires

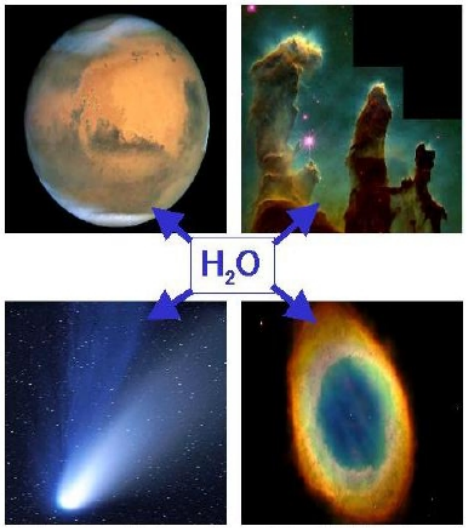

Alexandre Faure, Pierre Valiron, Michael Wernli et Laurent Wiesenfeld
Laboratoire d'Astrophysique de Grenoble

Thème : calcul des taux de désexcitation vibrationnelle du pliage de l'eau
$\rightarrow$ calculs de chimie quantique

environ 8 ans de calcul

CiGri: User examples

Simulation d'echo radar

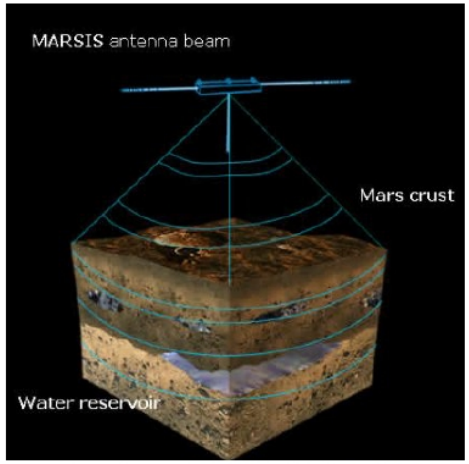

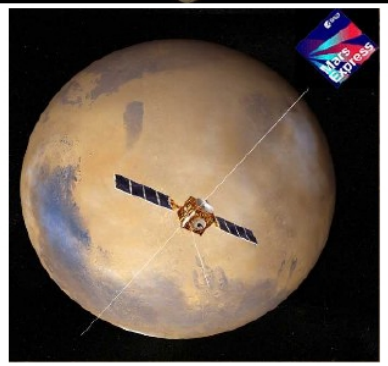

Jean-François Nouvel, Wlodek Kofman et Alain Herique
Laboratoire de Planétologie de Grenoble

Thème : Recherche de l'eau sur Mars par sondage radar

$\longrightarrow$ simulation pour optimiser la fréquence utilisée et pour sélectionner des orbites pertinentes (sonde Mars Express)

environ 719 jours de calcul

CiGri: User examples

Stratification de disques de poussières

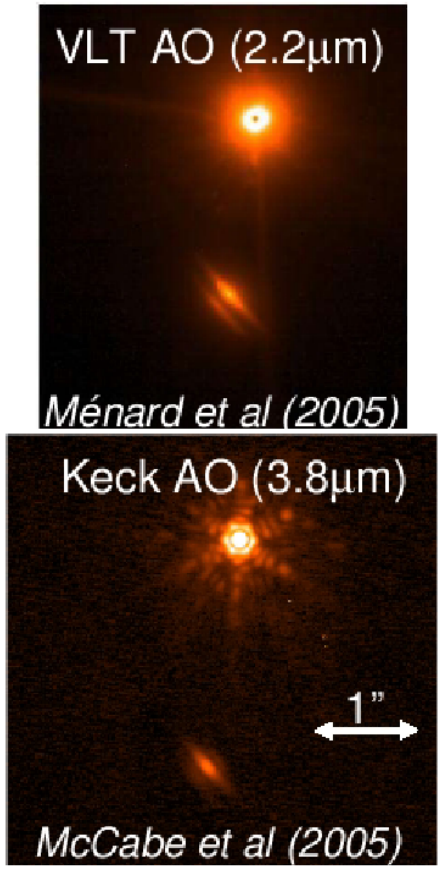

Christophe Pinte
Laboratoire d'Astrophysique de Grenoble

Thème : étude des disques de poussières
$\longrightarrow$ simulation de trajectoires de photons à travers des disques de poussières

environ 182 jours de calcul

CiGri: User examples

Astrocladistique : analyse phylogénétique de l'évolution des galaxies

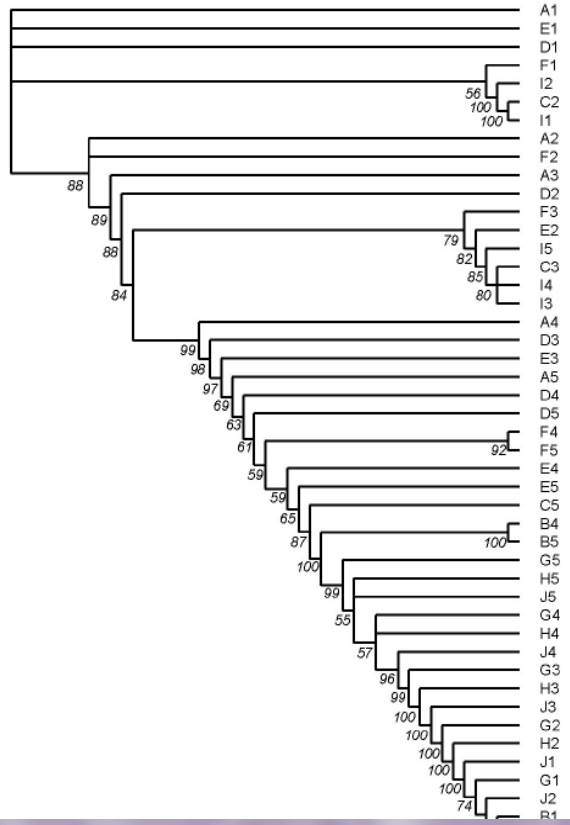

Didier Fraix-Burnet (LAOG, Grenoble), Emmanuel Davoust (LAOMP, Toulouse)

Thème : étude de la diversité des galaxies

$\longrightarrow$ organiser la diversité des galaxies selon leur état évolutif

683 jours de calcul

Conclusion

- CIMENT, c'est :
  - partage de machines + partage d'expertise
  - fédération structurée d'une communauté de calcul scientifique: des moyens, des scientifiques, des ingénieurs
  - des formations
  - des liens avec des projets scientifiques
- Un tel projet donne une unité a un mesocentre: de véritables moyens transversaux humains et matériels nécessaires

Questions?

ANNEXES

Cigri usage
1 - connect

- You need:
  - a login account on every cluster you want to use through the grid: contact the system administrator of those clusters, or your gridmaster if you know that he can do it for you.
  - a login account on the CiGri host: contact your gridmaster
  - the gridmaster has to establish a binding between your grid account and all your clusters accounts.

CiGri usage
1 - connect (cont.)

- As soon as you've got your grid account, you can connect on the web server of the cigri host where you'll find more informations on how to use the grid, status informations, usage statistics,...
- Also try to connect with SSH to the CiGri host
- Well, if you can connect, we can go further...

CiGri usage
2: validate your application

- Transfert your application on each cluster
- Compile it for the architecture of the cluster
- If it needs some input data, copy them
- Test your application with several parameters if possible
- Note*: to connect to the different clusters with **ssh**, you may need to connect to the CiGri host before, and use it as a gateway: some clusters are behind firewalls that do not accept connections from unknown hosts.*
You may need to generate an ssh key (see into the next section) for more comfort

CiGri usage
2: Validate your application (cont.)

Be carefull with NFS!

Si vous avez de gros volumes de données...

Une centaine de jobs qui essayent de charger (ou d'écrire) un fichier de plusieurs gigas situé sur votre home simultanément peuvent complètement saturer le serveur NFS.

Par ailleurs, une mauvaise optimisation du chargement de grosses données peut provoquer un overhead important anéantissant l'efficacité de votre campagne.

-> Utiliser les /scratch

-> Faire des seek au lieu de charger ou de parser complètement un fichier quand cela est possible

CiGri usage
3: campaign definition

- On the CiGri host, create a JDL (Job Description Language) file
  Example:

```
DEFAULT{
    name = campagne1 ;
    # njobs = 1000;
    paramFile = param.tmp;
}
icare.obs.ujf-grenoble.fr{
    execFile = /home/ciment/bzizou/test_Solaris10.sh ;
}
tomte.ujf-grenoble.fr{
    execFile = /users/nis/bzeznik/test.sh ;
}
```

CiGri usage
3: campaign definition (cont.)

- In the DEFAULT section, 2 modes:

```
paramFile = param.tmp;
  # Chaque ligne de ce fichier sera un argument de
  # votre programme (donc une tâche).

nbjobs = 1000;
  # L'application sera éxécutée 1000 fois.
```

- If using `paramFile` , the first application parameter will be used as the name of your job. Then, when CiGri will collect the results, stdout, stdin files and the file named like the first parameter will be copied to the CiGri host.

CiGri usage
3: campaign definition (cont.)

- You can ask for a data transfert (stage in) on submission, in the DEFAULT section:

```
data_to_transfer = /home/ciment/bzizou/demo/data/;
transfert_timeout = 20; # en secondes
```

- Warning: you need to generate an ssh key-pair on the CiGri host and install the public key on every cluster account. Here's an example of how you can do it:

Génération et diffusion d'une clé ssh

```
ssh-keygen -t rsa
for SERVER in one@idpot.imag.fr \
              elu@icare.obs.ujf-grenoble.fr \
              neo@tomte.ujf-grenoble.fr
do
  cat ~/.ssh/id_rsa.pub | \
     ssh $SERVER 'mkdir -p .ssh ; cat >> .ssh/authorized_keys'
done
```

CiGri usage
4: Submit the job

From the CiGri host:

gridsub -f nom_JDL

To cancel the entire campaign:

griddel -m <id>

To cancel a job:

griddel -j <id>

CiGri usage
5: Jobs visualization

- The web interface of the CiGri host shows the campaigns evolution (running jobs, executed jobs, waiting jobs, collected jobs,...) and statistics (jobs distributions over the clusters, time repartition, resubmission rate, estimated ending time,...)
- You can also watch errors and events. For example, if a cluster is unavailable, there's an event that can explain why and its status is « ToFix »

CiGri usage
6: Collect the results

- CiGri collects the results every 30 minutes
- The output files (stdout, stderr and the file whose name is the same as the first parameter) are backuped in a tgz (one tgz file per cluster)
- Results are put on the CiGri host into `~cigri/results/<username>/<campaign_number>/`

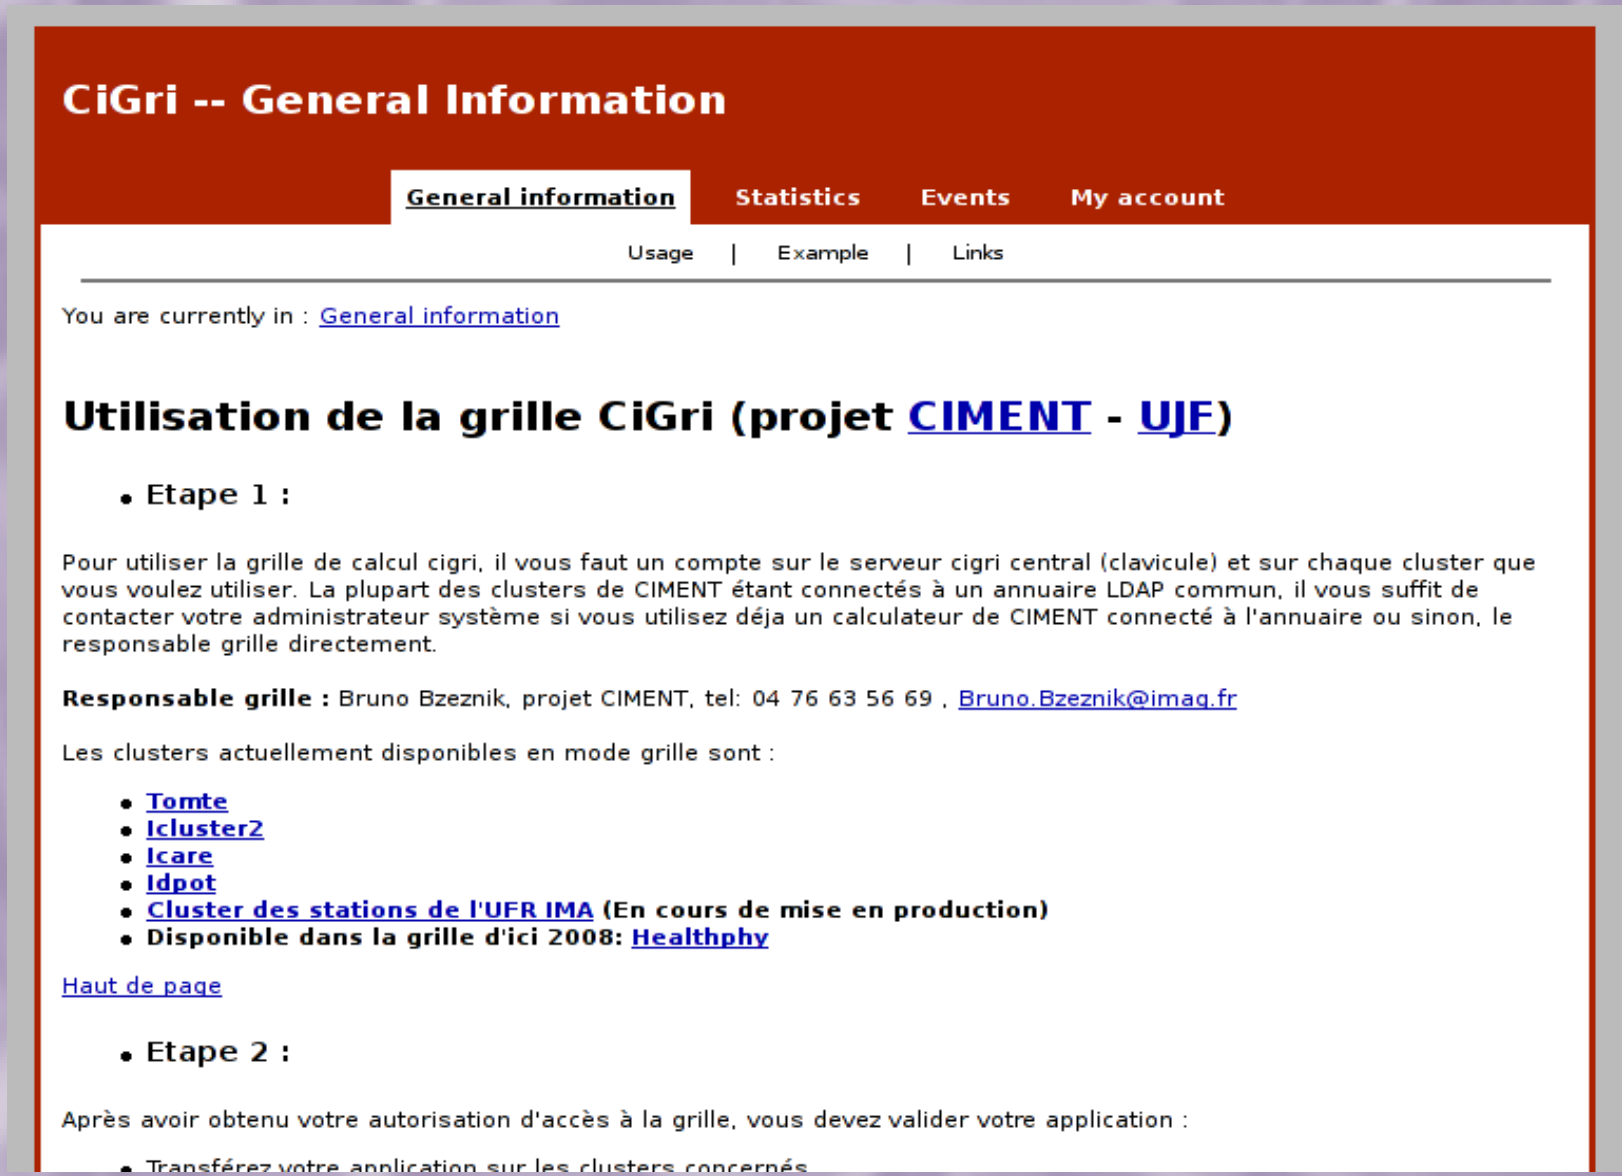
CiGri -- General Information
General information
Statistics
Events
My account
Usage | Example | Links
You are currently in : General information
Utilisation de la grille CiGri (projet CIMENT - UJF)
Etape 1 :
Pour utiliser la grille de calcul cigri, il vous faut un compte sur le serveur cigri central (clavicule) et sur chaque cluster que vous voulez utiliser. La plupart des clusters de CIMENT étant connectés à un annuaire LDAP commun, il vous suffit de contacter votre administrateur système si vous utilisez déja un calculateur de CIMENT connecté à l'annuaire ou sinon, le responsable grille directement.
Responsable grille : Bruno Bzeznik, projet CIMENT, tel: 04 76 63 56 69 , Bruno.Bzeznik@imag.fr
Les clusters actuellement disponibles en mode grille sont :
Tomte
Icluster2
Icare
Idpot
Cluster des stations de l'UFR IMA (En cours de mise en production)
Disponible dans la grille d'ici 2008: Healthphy
Haut de page
Etape 2 :
Après avoir obtenu votre autorisation d'accès à la grille, vous devez valider votre application :
Transférez votre application sur les clusters concernés

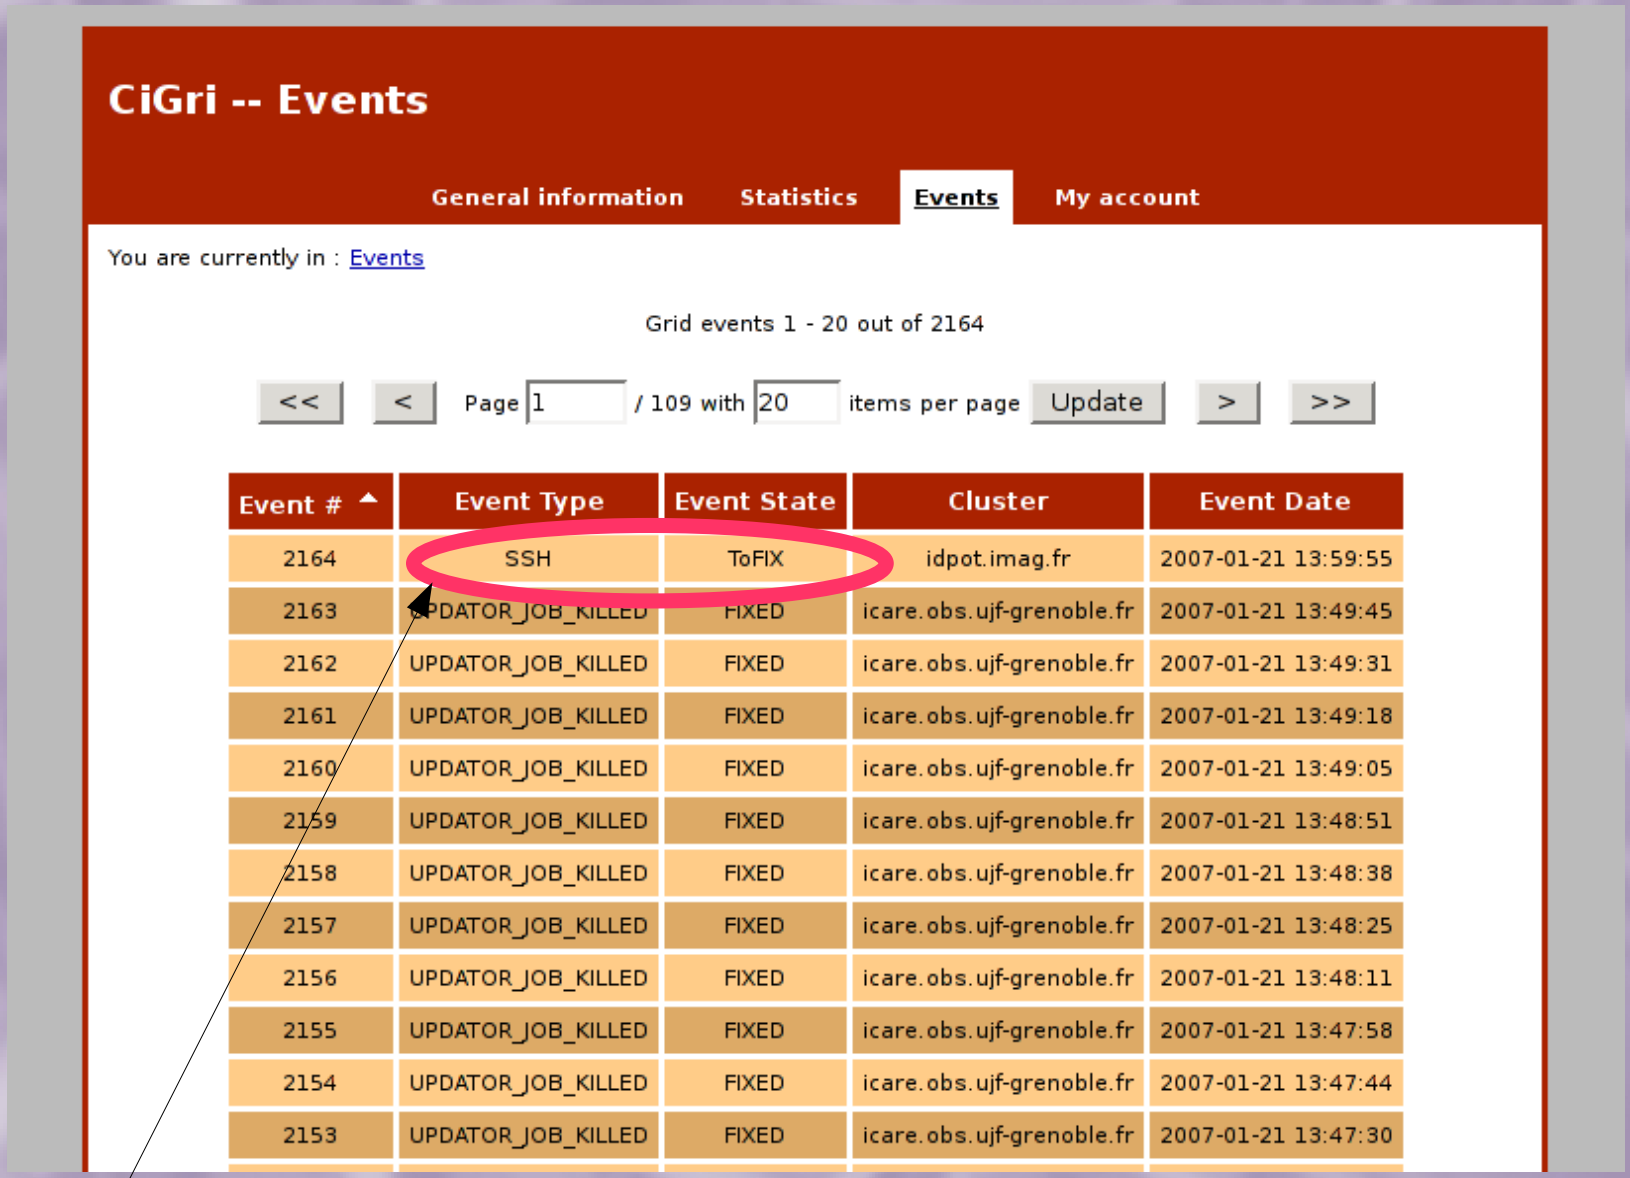

Contacter le gridmaster!

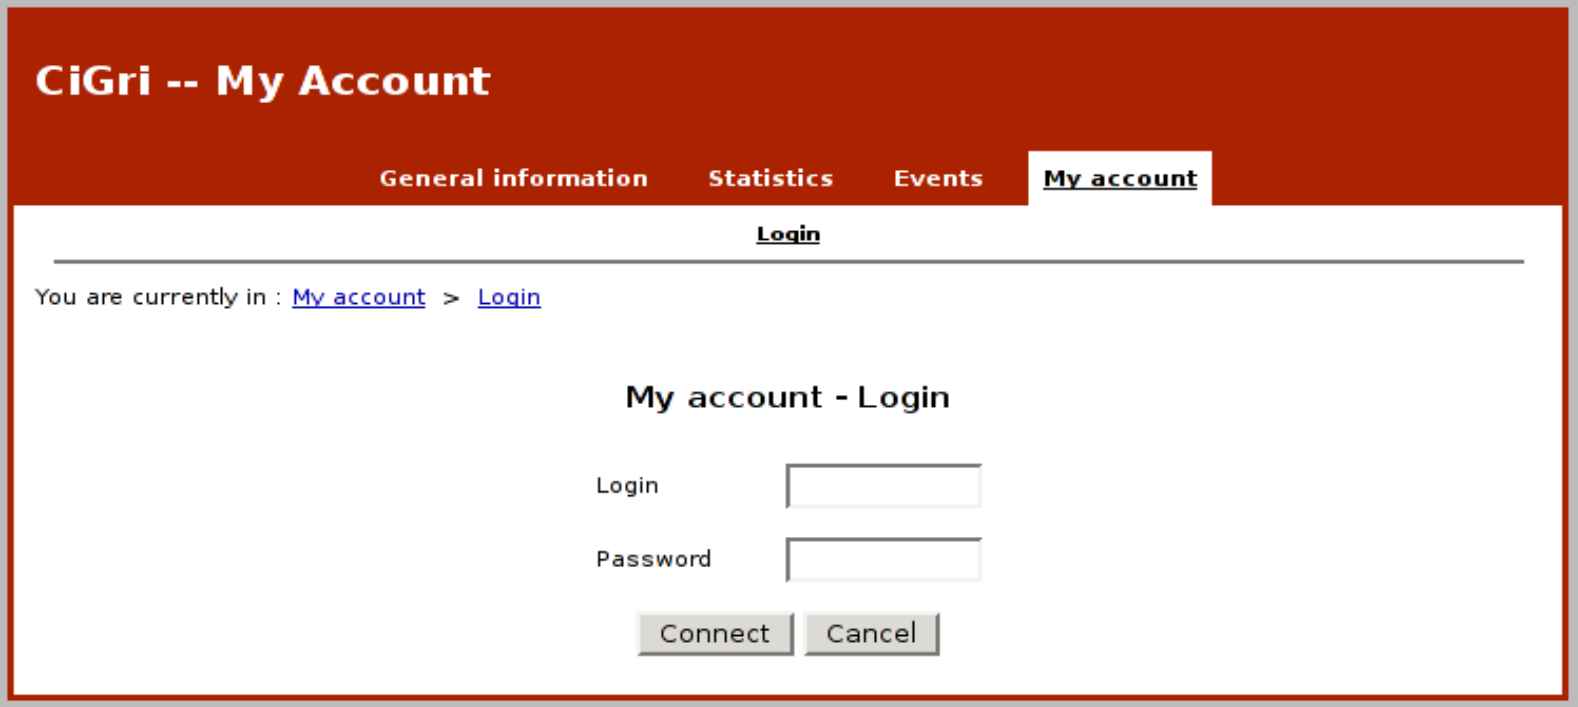
CiGri -- My Account
General information
Statistics
Events
My account
Login
You are currently in : My account > Login
My account - Login
Login
Password
Connect
Cancel
Contact: Bruno.Bzeznik@imag.fr

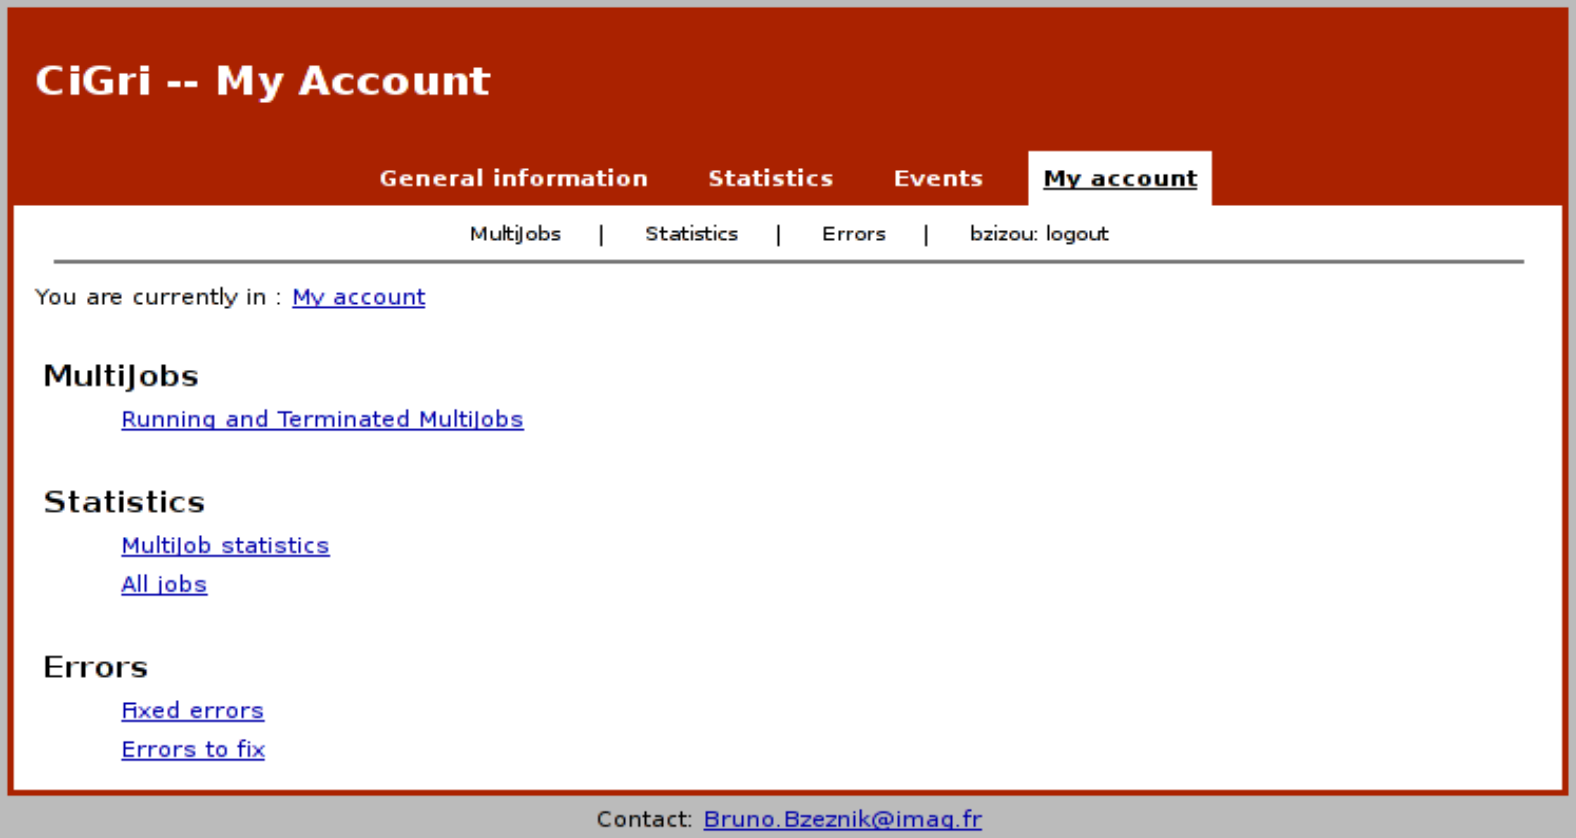
CiGri -- My Account
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account
MultiJobs
Running and Terminated MultiJobs
Statistics
MultiJob statistics
All jobs
Errors
Fixed errors
Errors to fix
Contact: Bruno.Bzeznik@imag.fr

The POVRAY example

- Povray: Persistence Of Vision is a raytracing software
- Eats a lot of CPU power

The POVRAY example

- Let's make a high definition animation of a virtual landscape by rotating a virtual camera.
- Source *landscape_rotate.pov*:

```
camera {
   location    <3.0, 2.3, 0.3+0.01*clock>
   direction   y
   sky         z
   up          z
   right       (4/3)*x
   look_at     <0.0, 0.0, 0.1>
   angle       32
   rotate <0,0,clock>
}
```

The POVRAY example

- I compiled povray 3.6 in my home directory of every clusters into ~/povray-3.6
- A frame is computed, for a given camera angle, like that:

```
povray +W1280 +H768 +L -K180 +Olandscape0180.png \
                                        landscape_camera.pov
```

- An entire circle: 360 degrees, so 360 jobs with the following changing parameters
  - The name of the output file
  - The camera angle

The POVRAY example

• You have to make a script (on clusters) witch takes parameters into the good order (remember that CiGri uses the first parameter as the output filename):

./run_povray.sh <output_file> <input_file> <x_size> <y_size> <clock>

```
#!/bin/bash
RUN_DIR="${HOME}/povray"
POV_INC=povray-3.6/include
POV_CMD=povray-3.6/povray
# Args
if [ $# -ne 5 ]; then
    echo "ERROR: Bad number of parameters"
    echo "USAGE: $0 <output_file> <input_file> <x_size> <y_size>
<clock>"
    exit 1
fi
# Compute
cd $RUN_DIR
${POV_CMD} +W${3} +H${4} +L${POV_INC} -K$5 +O${1} ${2}
```

The POVRAY example

- Now we can generate the parameters (bash example):

```
n=0;for i in `seq 0 2 360`; do printf "landscape_
%0.4d.png landscape_rotate.pov 1280 768 $i\n" $n;
let n++;done > params.tmp
```

```
# Wich gives us a file called params.tmp:
landscape_0000.png landscape_rotate.pov 1280 768 0
landscape_0001.png landscape_rotate.pov 1280 768 2
landscape_0002.png landscape_rotate.pov 1280 768 4
landscape_0003.png landscape_rotate.pov 1280 768 6
landscape_0004.png landscape_rotate.pov 1280 768 8
landscape_0005.png landscape_rotate.pov 1280 768 10
landscape_0006.png landscape_rotate.pov 1280 768 12
landscape_0007.png landscape_rotate.pov 1280 768 14
landscape_0008.png landscape_rotate.pov 1280 768 16
landscape_0009.png landscape_rotate.pov 1280 768 18
...
```

The POVRAY example

- And the job description file (JDL):

```
DEFAULT{
    name = pov_ciment_demo ;
    paramFile = params.tmp;
}
tomte.ujf-grenoble.fr{
    execFile = /home/nis/bzeznik/povray/run_povray.sh ;
    walltime = 00:15:00 ;
    weight = 1;
    execDir = /home/nis/bzeznik/povray/ ;
}
icare.obs.ujf-grenoble.fr{
    execFile = /users/bzeznik/povray/run_povray_solaris.sh;
    walltime = 00:10:00 ;
    weight = 1;
    execDir = /users/bzeznik/povray/ ;
}
idpot.imag.fr{
    execFile = /home/grenoble/bbzeznik/povray/run_povray.sh ;
    walltime = 00:10:00 ;
    weight = 1;
    execDir = /home/grenoble/bbzeznik/povray/ ;
}
```

The POVRAY example

Submission:

```
bzizou@clavicule:~/demo_povray$ gridsub -f pov_camera_JDL
```

```
DL file = pov_camera_JDL
nsert (29,'landscape_0000.png ./data/landscape_rotate.pov 1280 768 0','landscape_0000.png')
nsert (29,'landscape_0001.png ./data/landscape_rotate.pov 1280 768 2','landscape_0001.png')
nsert (29,'landscape_0002.png ./data/landscape_rotate.pov 1280 768 4','landscape_0002.png')
nsert (29,'landscape_0003.png ./data/landscape_rotate.pov 1280 768 6','landscape_0003.png')
nsert (29,'landscape_0004.png ./data/landscape_rotate.pov 1280 768 8','landscape_0004.png')
nsert (29,'landscape_0005.png ./data/landscape_rotate.pov 1280 768 10','landscape_0005.png')
nsert (29,'landscape_0006.png ./data/landscape_rotate.pov 1280 768 12','landscape_0006.png')
nsert (29,'landscape_0007.png ./data/landscape_rotate.pov 1280 768 14','landscape_0007.png')
nsert (29,'landscape_0008.png ./data/landscape_rotate.pov 1280 768 16','landscape_0008.png')
nsert (29,'landscape_0009.png ./data/landscape_rotate.pov 1280 768 18','landscape_0009.png')
nsert (29,'landscape_0010.png ./data/landscape_rotate.pov 1280 768 20','landscape_0010.png')
nsert (29,'landscape_0011.png ./data/landscape_rotate.pov 1280 768 22','landscape_0011.png')
nsert (29,'landscape_0012.png ./data/landscape_rotate.pov 1280 768 24','landscape_0012.png')
nsert (29,'landscape_0013.png ./data/landscape_rotate.pov 1280 768 26','landscape_0013.png')
nsert (29,'landscape_0014.png ./data/landscape_rotate.pov 1280 768 28','landscape_0014.png')
nsert (29,'landscape_0015.png ./data/landscape_rotate.pov 1280 768 30','landscape_0015.png')
nsert (29,'landscape_0016.png ./data/landscape_rotate.pov 1280 768 32','landscape_0016.png')
nsert (29,'landscape_0017.png ./data/landscape_rotate.pov 1280 768 34','landscape_0017.png')
..
```

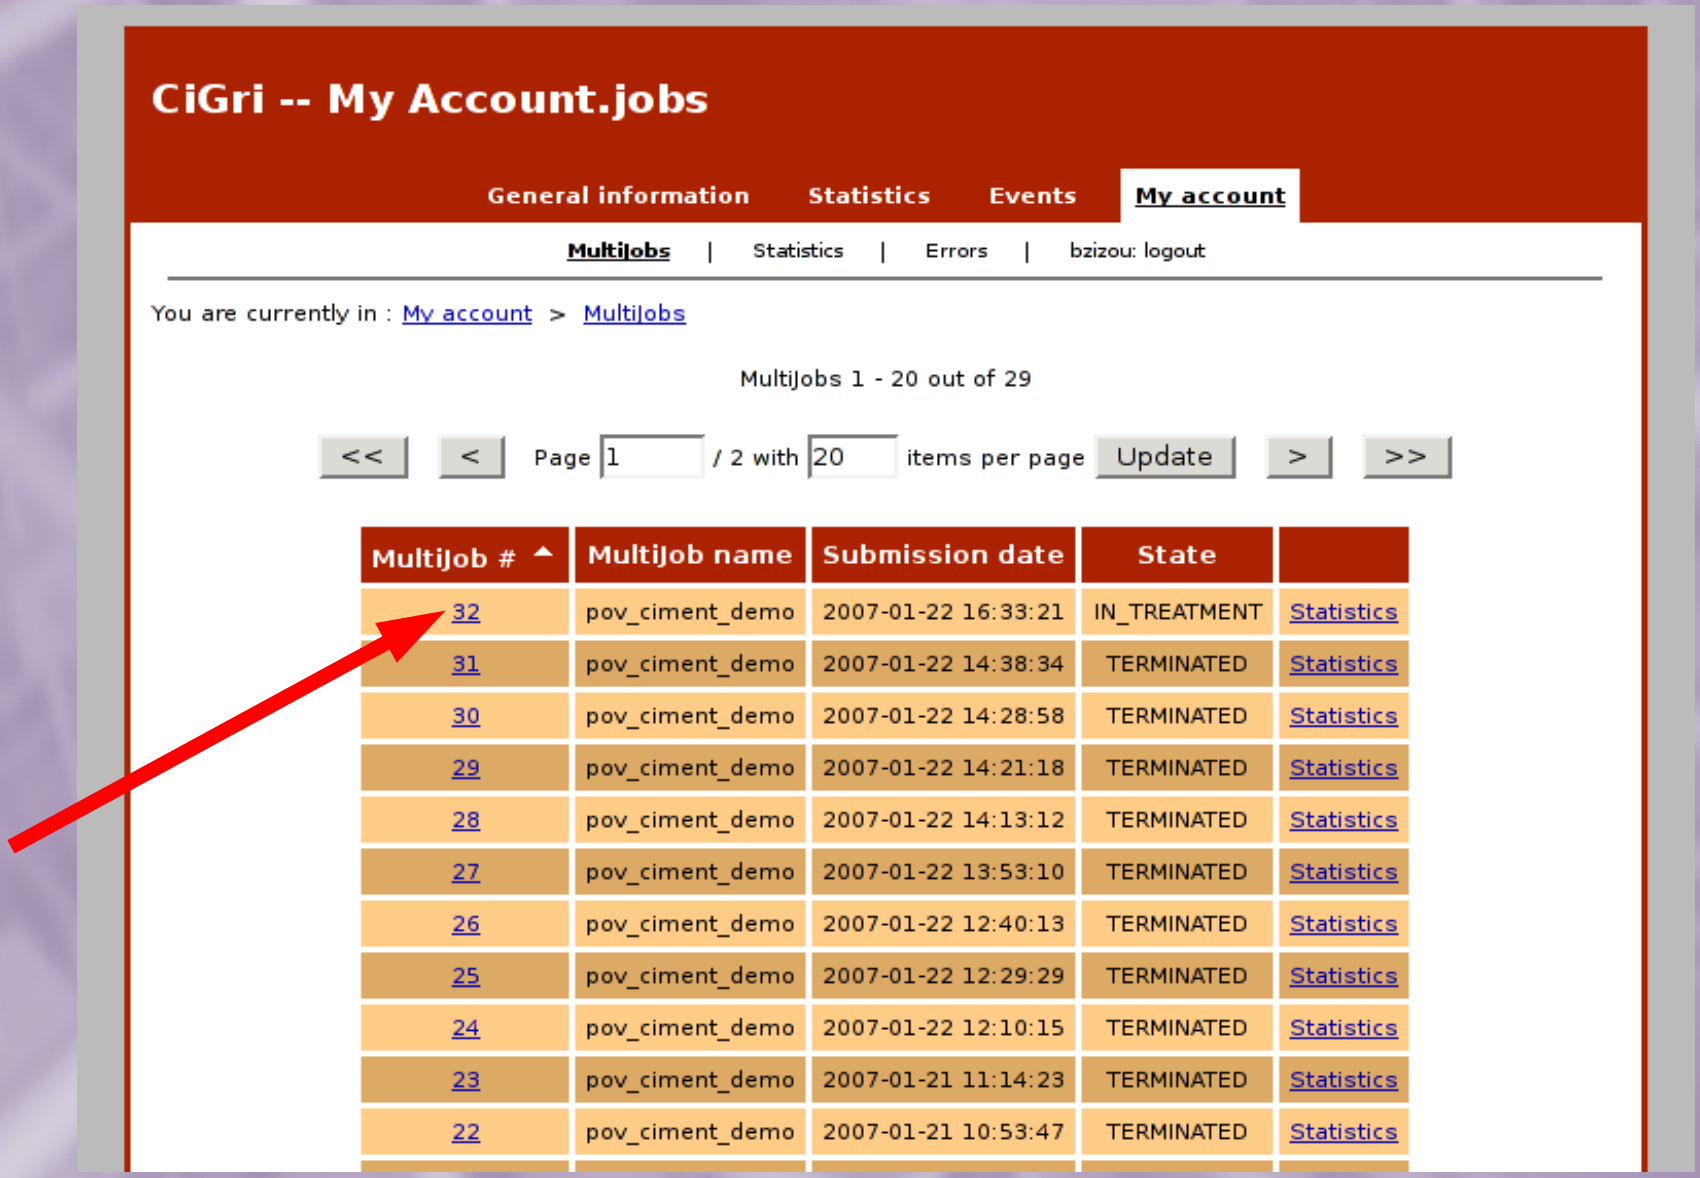
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs
MultiJobs 1 - 20 out of 29
<< < Page 1 / 2 with 20 items per page Update > >>
MultiJob # | MultiJob name | Submission date | State |
32 | pov_ciment_demo | 2007-01-22 16:33:21 | IN_TREATMENT | Statistics
31 | pov_ciment_demo | 2007-01-22 14:38:34 | TERMINATED | Statistics
30 | pov_ciment_demo | 2007-01-22 14:28:58 | TERMINATED | Statistics
29 | pov_ciment_demo | 2007-01-22 14:21:18 | TERMINATED | Statistics
28 | pov_ciment_demo | 2007-01-22 14:13:12 | TERMINATED | Statistics
27 | pov_ciment_demo | 2007-01-22 13:53:10 | TERMINATED | Statistics
26 | pov_ciment_demo | 2007-01-22 12:40:13 | TERMINATED | Statistics
25 | pov_ciment_demo | 2007-01-22 12:29:29 | TERMINATED | Statistics
24 | pov_ciment_demo | 2007-01-22 12:10:15 | TERMINATED | Statistics
23 | pov_ciment_demo | 2007-01-21 11:14:23 | TERMINATED | Statistics
22 | pov_ciment_demo | 2007-01-21 10:53:47 | TERMINATED | Statistics

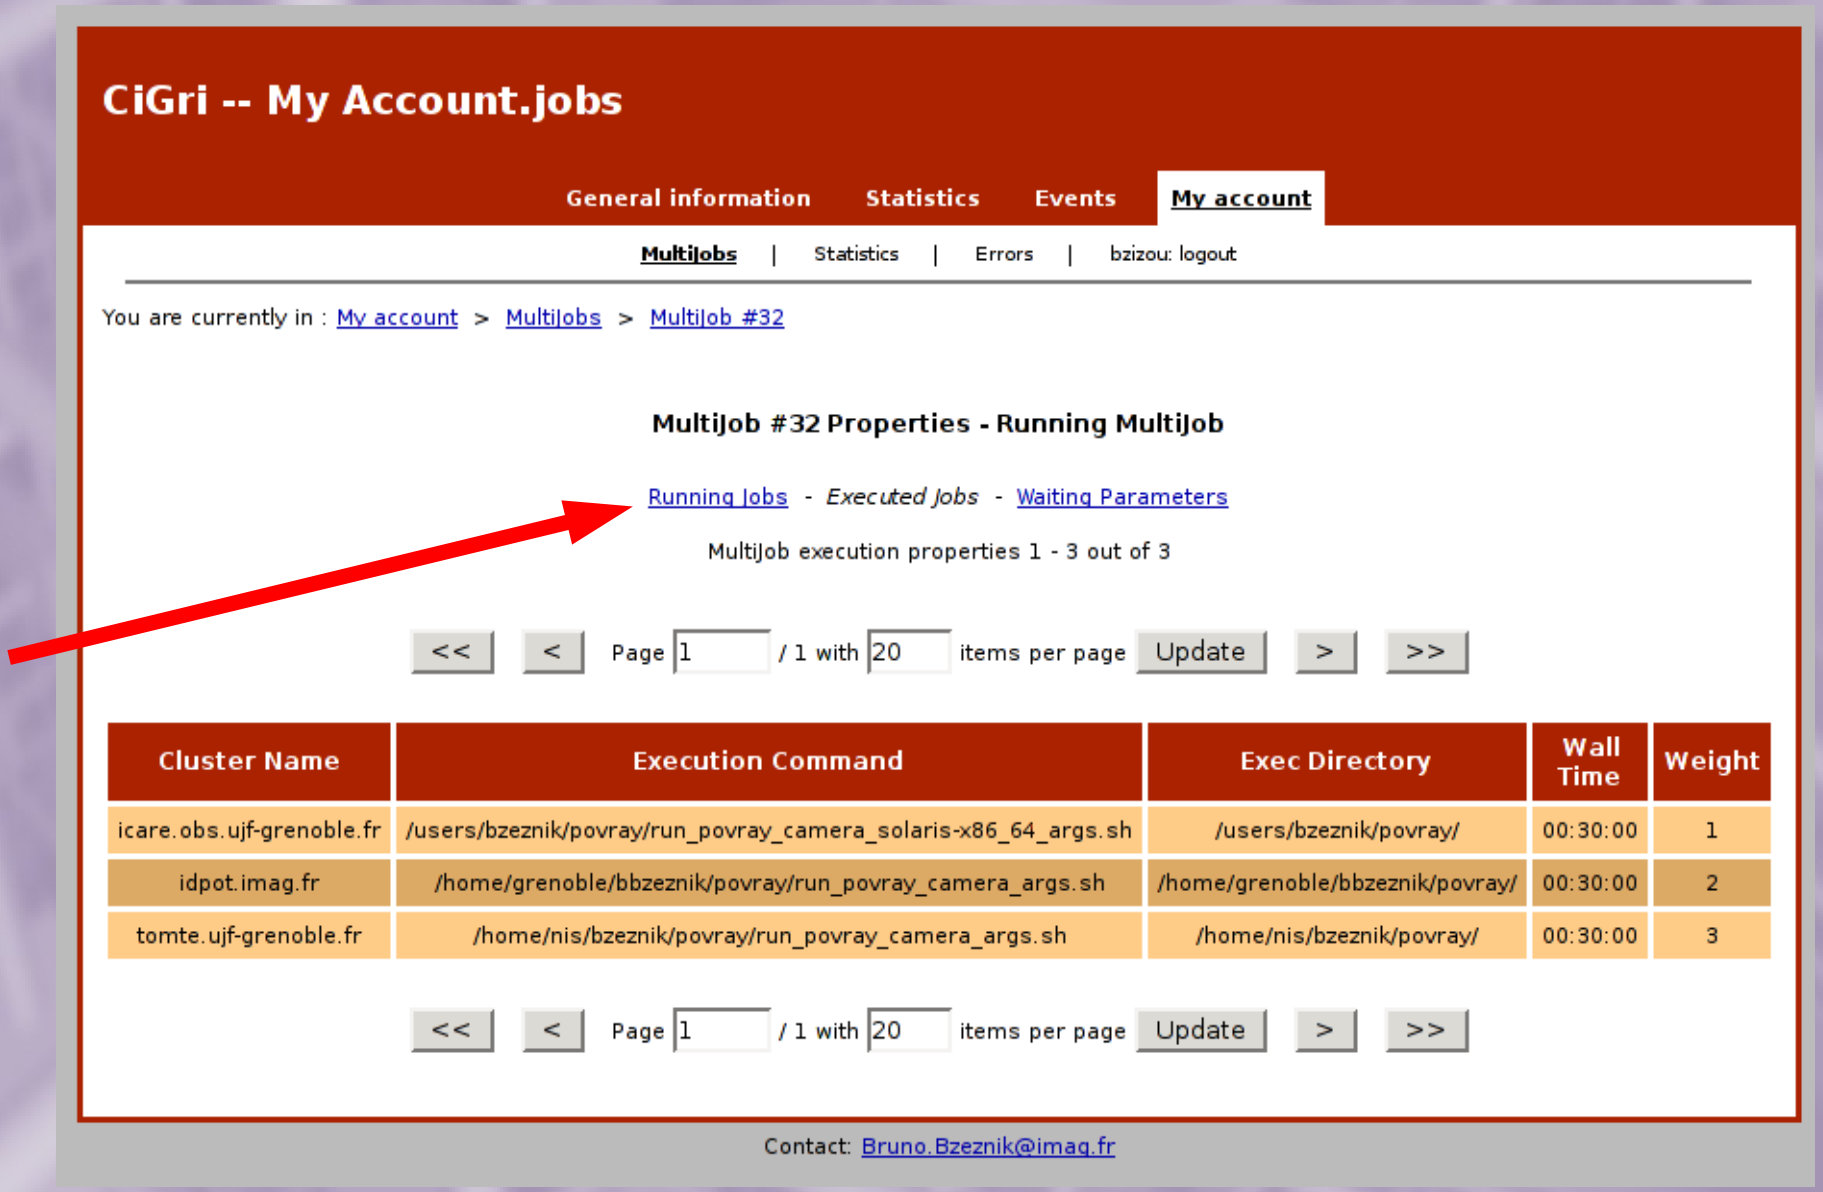
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32
MultiJob #32 Properties - Running MultiJob
Running Jobs - Executed Jobs - Waiting Parameters
MultiJob execution properties 1 - 3 out of 3
<< < Page 1 / 1 with 20 items per page Update > >>
Cluster Name
Execution Command
Exec Directory
Wall Time
Weight
icare.obs.ujf-grenoble.fr
/users/bzeznik/povray/run_povray_camera_solaris-x86_64_args.sh
/users/bzeznik/povray/
00:30:00
1
idpot.imag.fr
/home/grenoble/bbzeznik/povray/run_povray_camera_args.sh
/home/grenoble/bbzeznik/povray/
00:30:00
2
tomte.ujf-grenoble.fr
/home/nis/bzeznik/povray/run_povray_camera_args.sh
/home/nis/bzeznik/povray/
00:30:00
3
<< < Page 1 / 1 with 20 items per page Update > >>
Contact: Bruno.Bzeznik@imag.fr

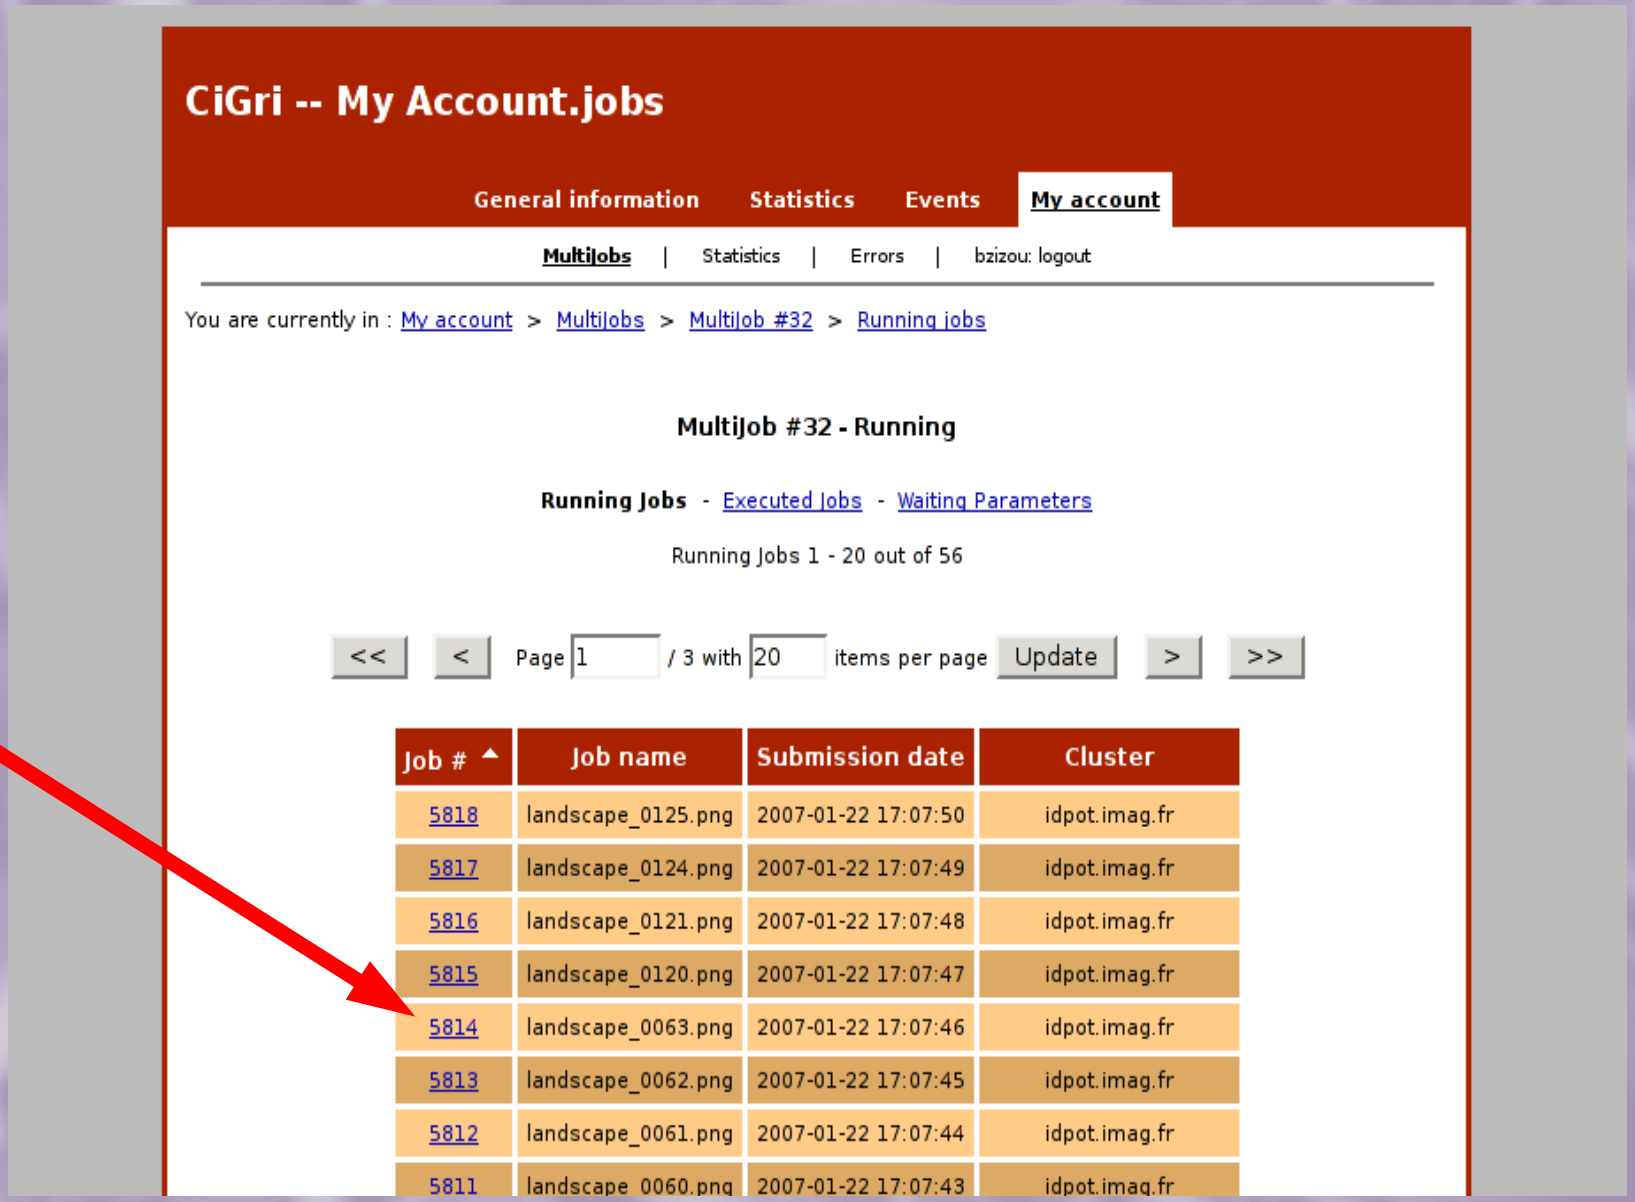
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32 > Running jobs
MultiJob #32 - Running
Running Jobs - Executed Jobs - Waiting Parameters
Running Jobs 1 - 20 out of 56
<< < Page 1 / 3 with 20 items per page Update > >>
Job # | Job name | Submission date | Cluster
5818 | landscape_0125.png | 2007-01-22 17:07:50 | idpot.imag.fr
5817 | landscape_0124.png | 2007-01-22 17:07:49 | idpot.imag.fr
5816 | landscape_0121.png | 2007-01-22 17:07:48 | idpot.imag.fr
5815 | landscape_0120.png | 2007-01-22 17:07:47 | idpot.imag.fr
5814 | landscape_0063.png | 2007-01-22 17:07:46 | idpot.imag.fr
5813 | landscape_0062.png | 2007-01-22 17:07:45 | idpot.imag.fr
5812 | landscape_0061.png | 2007-01-22 17:07:44 | idpot.imag.fr
5811 | landscape_0060.png | 2007-01-22 17:07:43 | idpot.imag.fr


CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32 > Running jobs > Job #5814 details
Job #5814 in MultiJob #32
Job Name
landscape_0063.png
Job Parameters
landscape_0063.png landscape_rotate.pov 1280 768 63
Job State
Running
Cluster Name
idpot.imag.fr
Node Name
Submission date
2007-01-22 17:07:46
Start date
End date
Duration
00:00:00
Collect #
0
Batch Id
1836
Return Code
Contact: Bruno.Bzeznik@imag.fr

```
bbzeznik@idpot:~/povray$ tail OAR.1836.stderr
  0:00:05 Creating bounding slabs
  0:00:05 Creating vista buffer
  0:00:05 Creating light buffers 269K tokens
Scene Statistics
  Finite objects:          3330
  Infinite objects:           0
  Light sources:              1
  Total:                   3331

  0:05:07 Rendering line 139 of 768
```

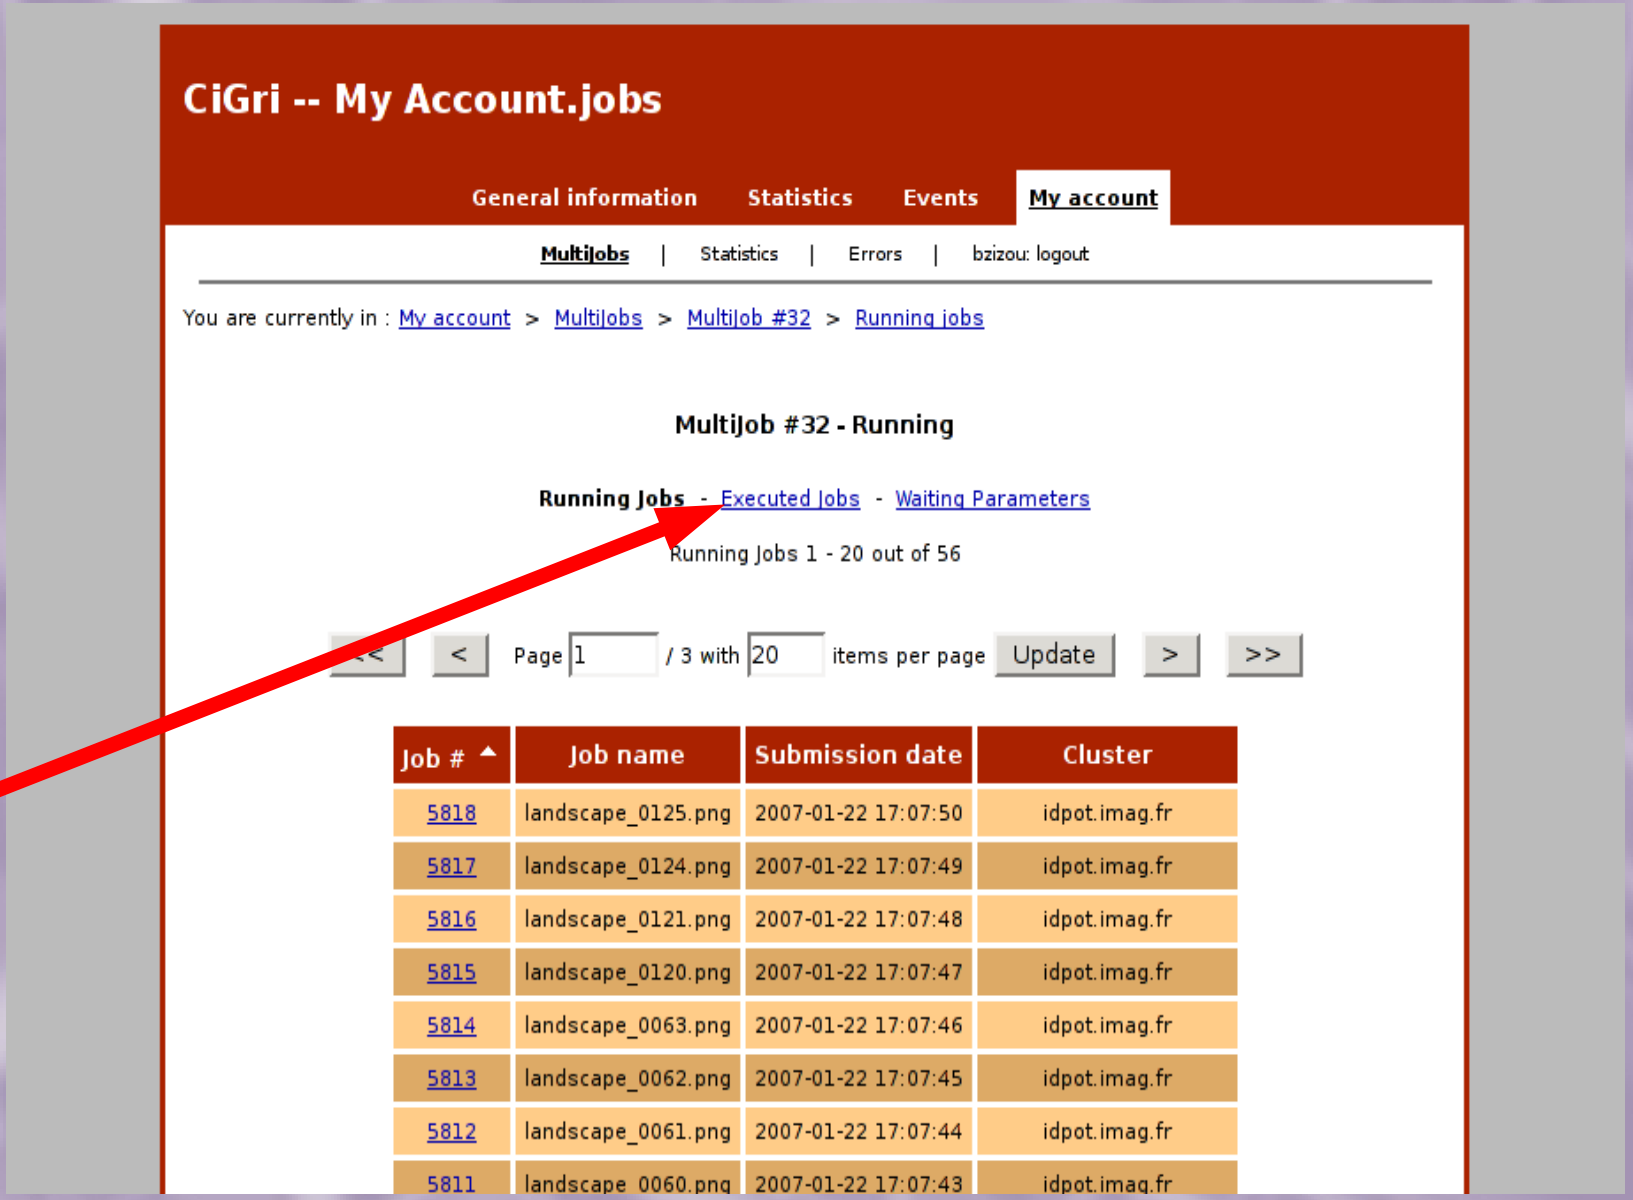
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32 > Running jobs
MultiJob #32 - Running
Running Jobs - Executed Jobs - Waiting Parameters
Running Jobs 1 - 20 out of 56
<< < Page 1 / 3 with 20 items per page Update > >>
Job # | Job name | Submission date | Cluster
5818 | landscape_0125.png | 2007-01-22 17:07:50 | idpot.imag.fr
5817 | landscape_0124.png | 2007-01-22 17:07:49 | idpot.imag.fr
5816 | landscape_0121.png | 2007-01-22 17:07:48 | idpot.imag.fr
5815 | landscape_0120.png | 2007-01-22 17:07:47 | idpot.imag.fr
5814 | landscape_0063.png | 2007-01-22 17:07:46 | idpot.imag.fr
5813 | landscape_0062.png | 2007-01-22 17:07:45 | idpot.imag.fr
5812 | landscape_0061.png | 2007-01-22 17:07:44 | idpot.imag.fr
5811 | landscape_0060.png | 2007-01-22 17:07:43 | idpot.imag.fr

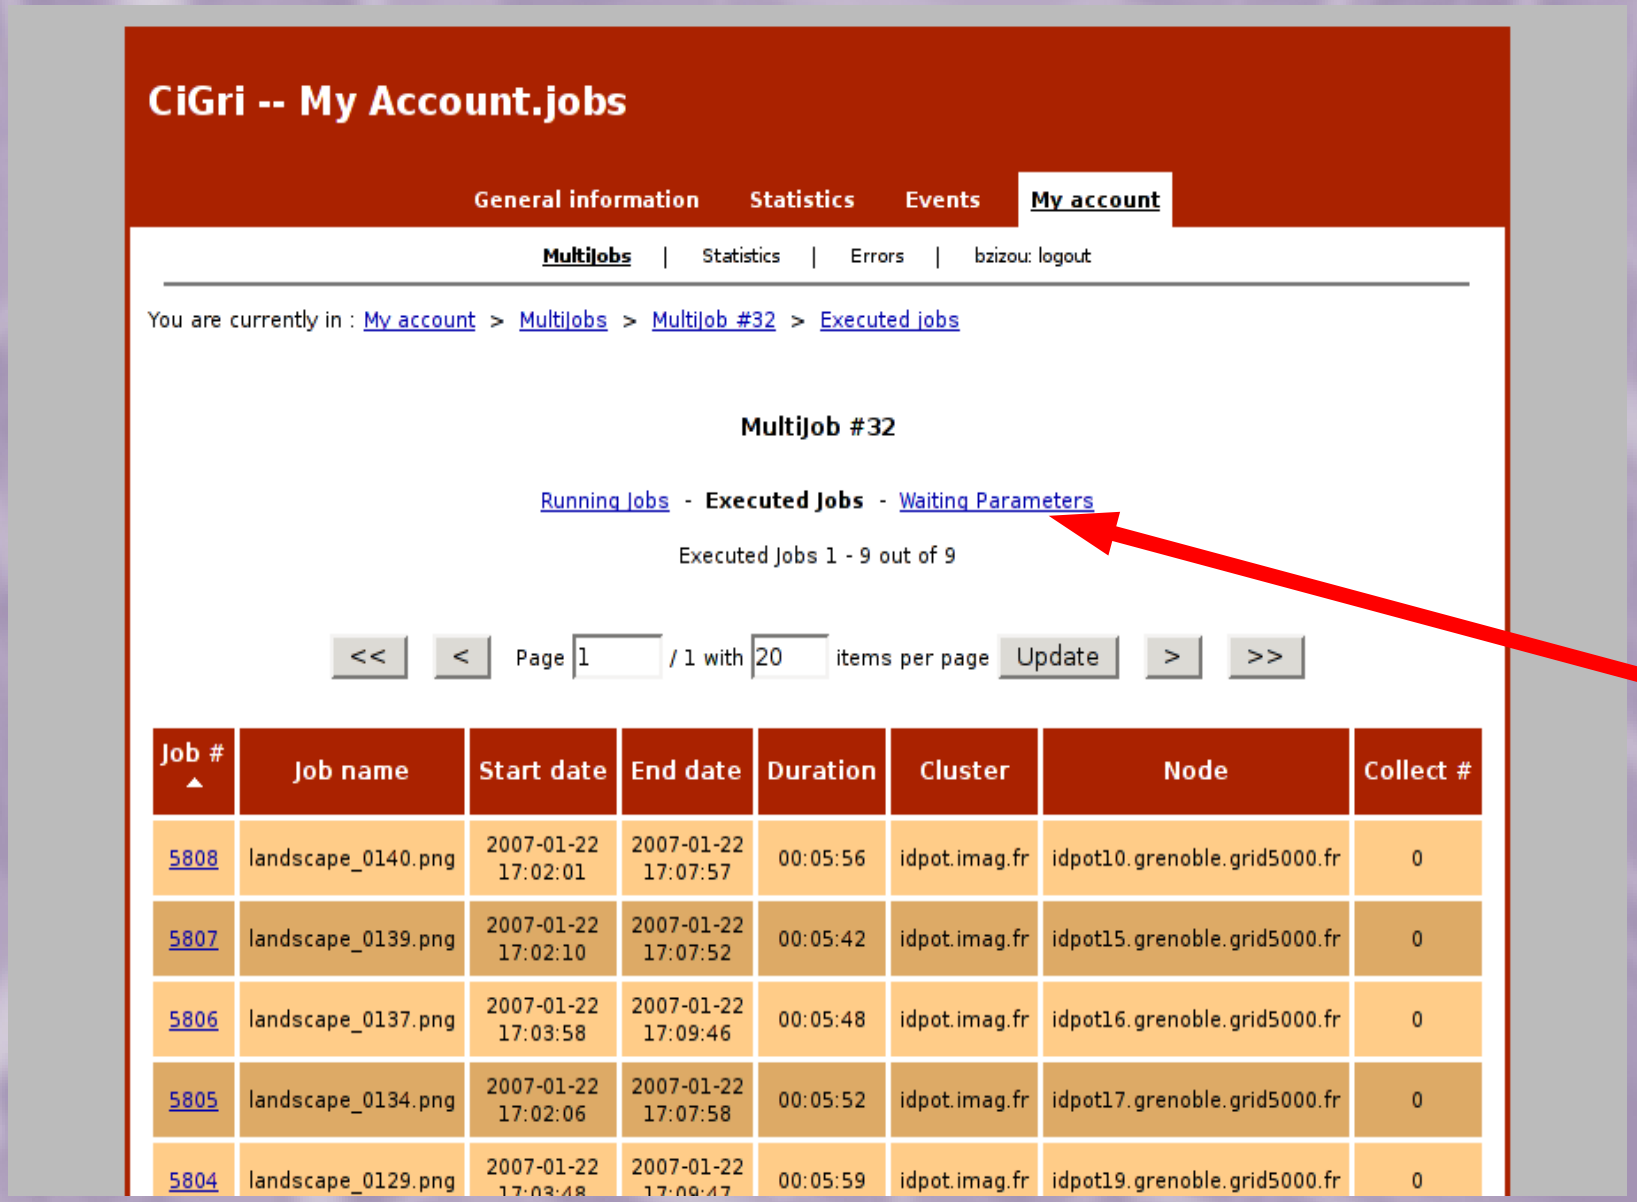
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32 > Executed jobs
MultiJob #32
Running Jobs - Executed Jobs - Waiting Parameters
Executed Jobs 1 - 9 out of 9
<< < Page 1 / 1 with 20 items per page Update > >>
Job # | Job name | Start date | End date | Duration | Cluster | Node | Collect #
5808 | landscape_0140.png | 2007-01-22 17:02:01 | 2007-01-22 17:07:57 | 00:05:56 | idpot.imag.fr | idpot10.grenoble.grid5000.fr | 0
5807 | landscape_0139.png | 2007-01-22 17:02:10 | 2007-01-22 17:07:52 | 00:05:42 | idpot.imag.fr | idpot15.grenoble.grid5000.fr | 0
5806 | landscape_0137.png | 2007-01-22 17:03:58 | 2007-01-22 17:09:46 | 00:05:48 | idpot.imag.fr | idpot16.grenoble.grid5000.fr | 0
5805 | landscape_0134.png | 2007-01-22 17:02:06 | 2007-01-22 17:07:58 | 00:05:52 | idpot.imag.fr | idpot17.grenoble.grid5000.fr | 0
5804 | landscape_0129.png | 2007-01-22 | 2007-01-22 | 00:05:59 | idpot.imag.fr | idpot19.grenoble.grid5000.fr | 0

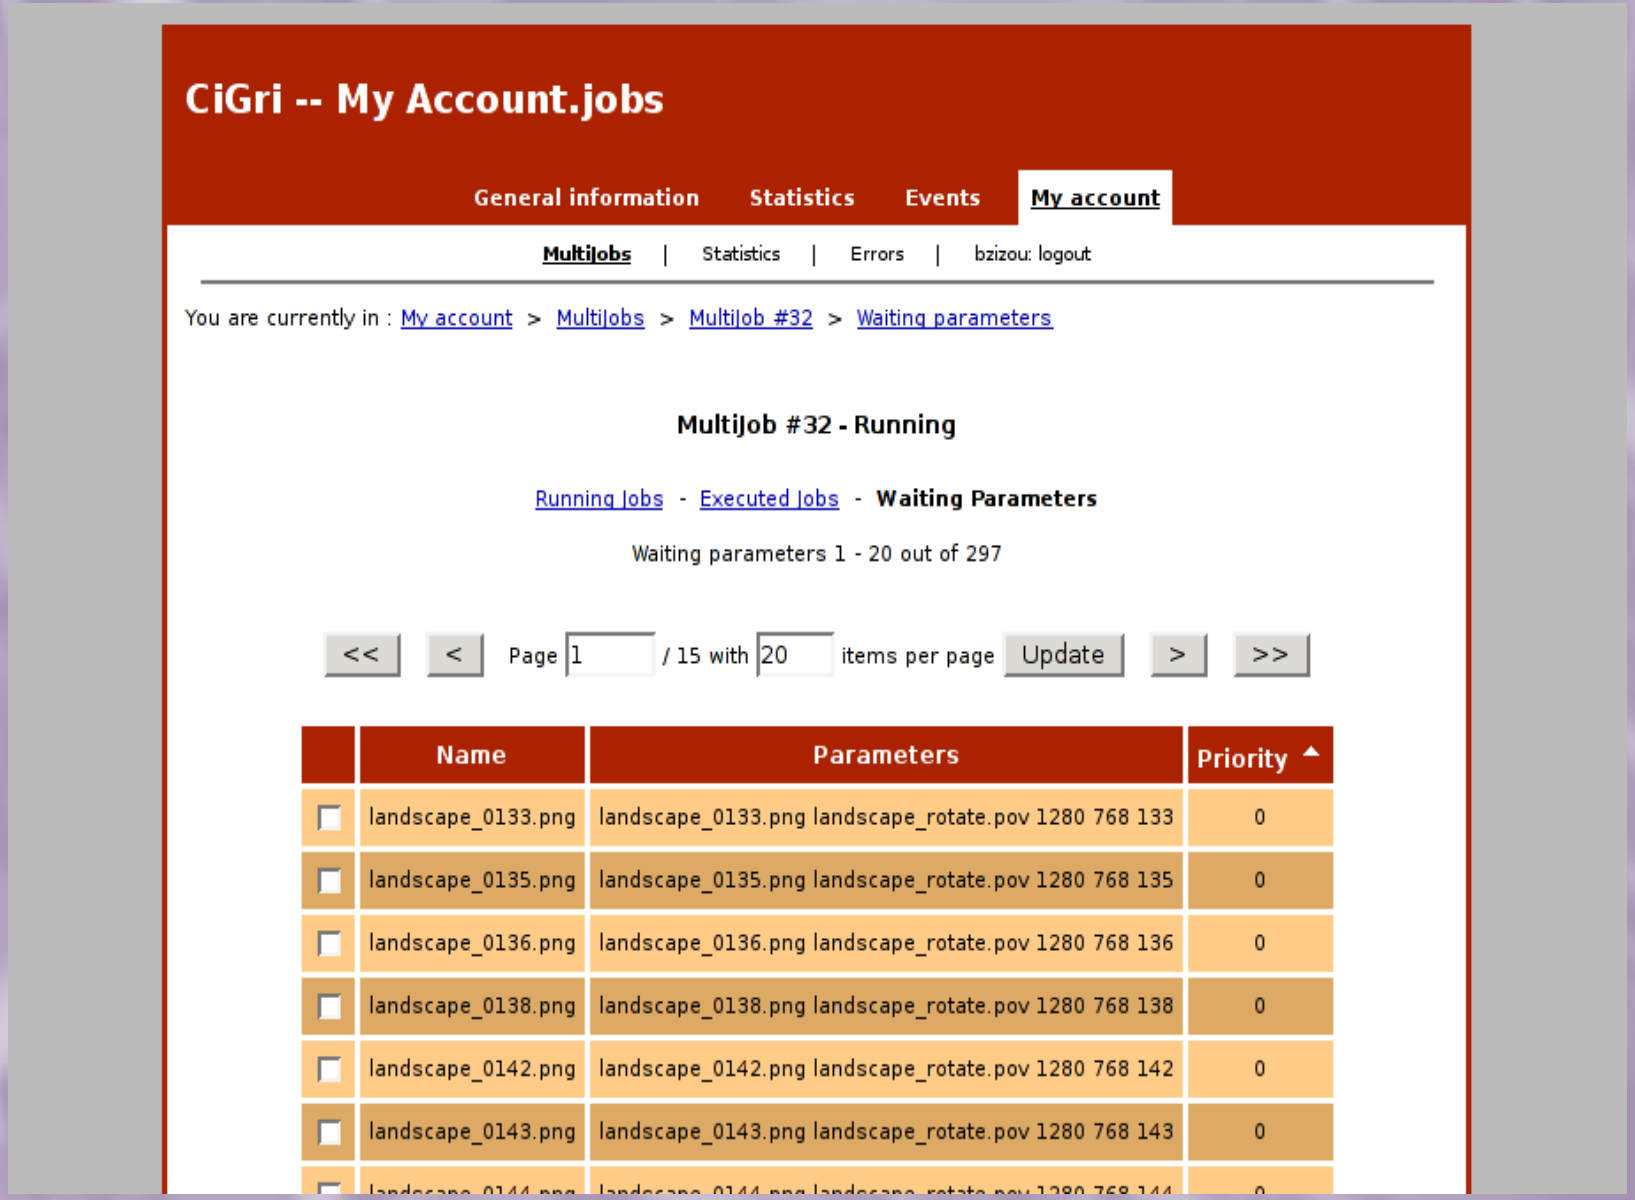
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs > MultiJob #32 > Waiting parameters
MultiJob #32 - Running
Running Jobs - Executed Jobs - Waiting Parameters
Waiting parameters 1 - 20 out of 297
<< < Page 1 / 15 with 20 items per page Update > >>
Name
Parameters
Priority
landscape_0133.png landscape_0133.png landscape_rotate.pov 1280 768 133 0
landscape_0135.png landscape_0135.png landscape_rotate.pov 1280 768 135 0
landscape_0136.png landscape_0136.png landscape_rotate.pov 1280 768 136 0
landscape_0138.png landscape_0138.png landscape_rotate.pov 1280 768 138 0
landscape_0142.png landscape_0142.png landscape_rotate.pov 1280 768 142 0
landscape_0143.png landscape_0143.png landscape_rotate.pov 1280 768 143 0

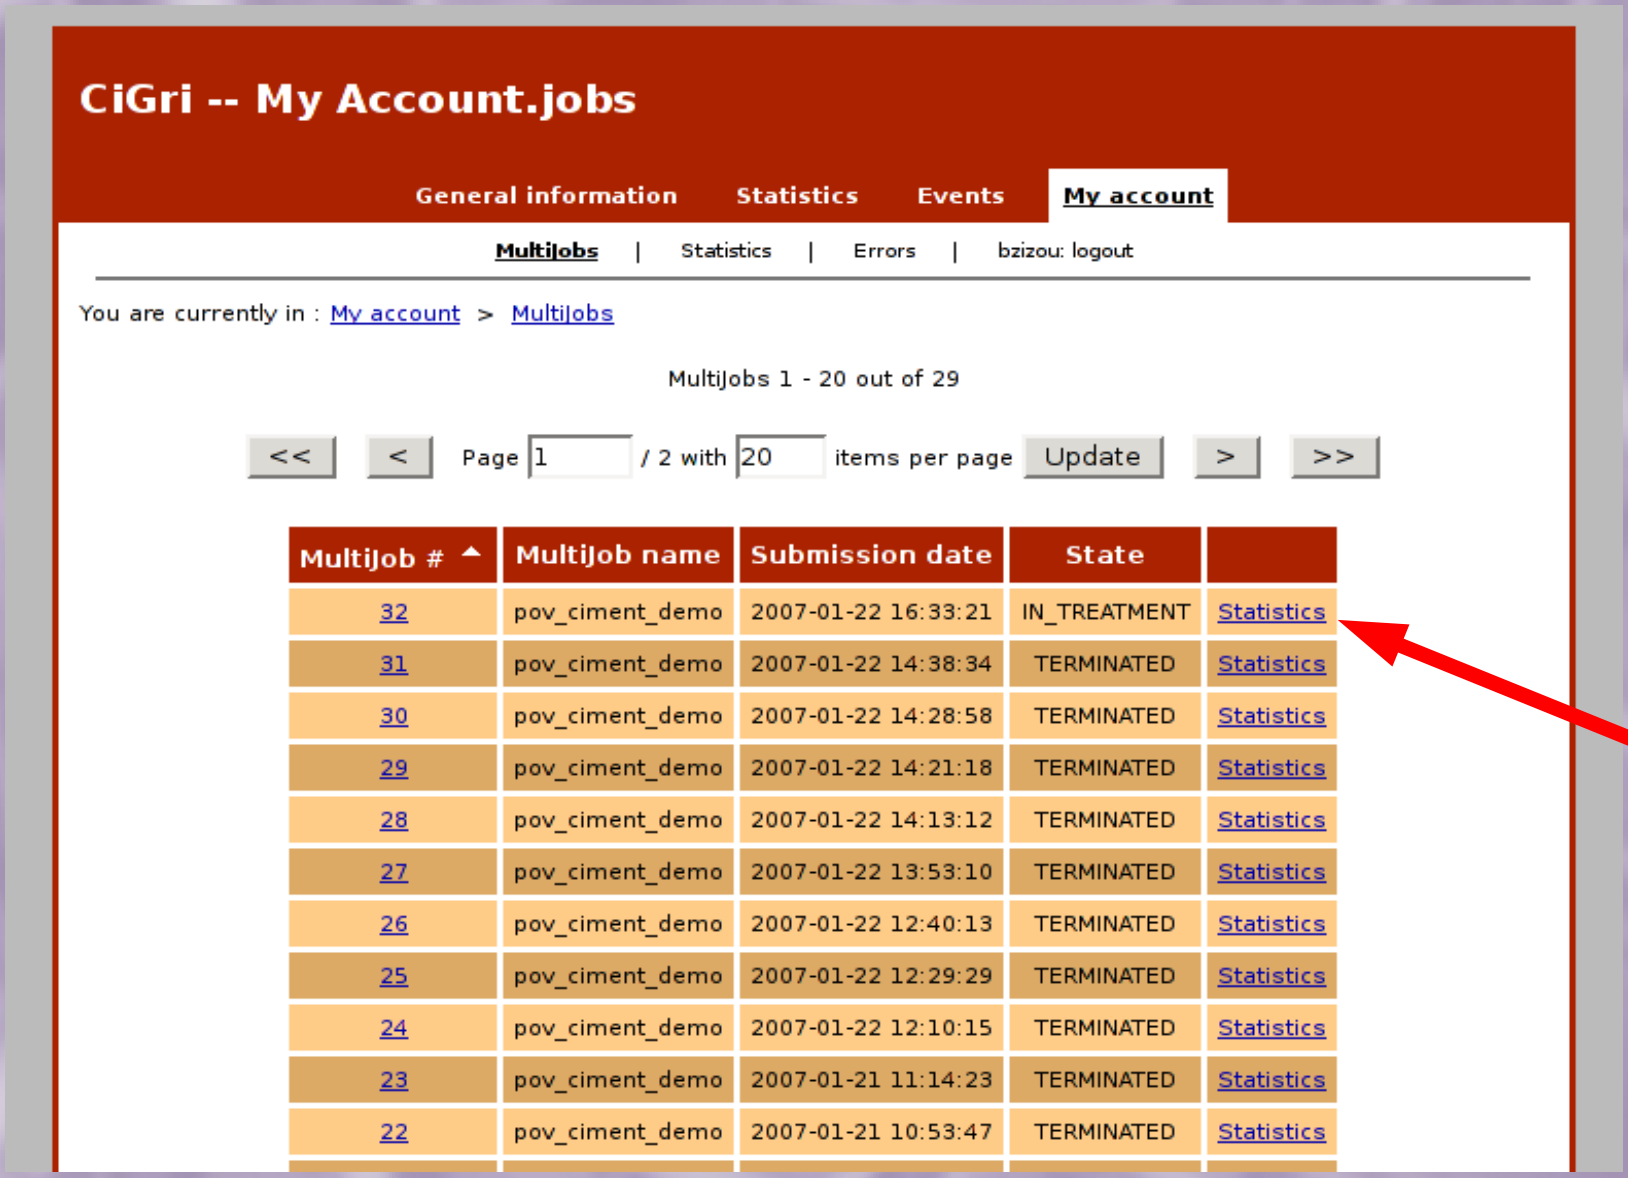
CiGri -- My Account.jobs
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > MultiJobs
MultiJobs 1 - 20 out of 29
<< < Page 1 / 2 with 20 items per page Update > >>
MultiJob # | MultiJob name | Submission date | State |
32 | pov_ciment_demo | 2007-01-22 16:33:21 | IN_TREATMENT | Statistics
31 | pov_ciment_demo | 2007-01-22 14:38:34 | TERMINATED | Statistics
30 | pov_ciment_demo | 2007-01-22 14:28:58 | TERMINATED | Statistics
29 | pov_ciment_demo | 2007-01-22 14:21:18 | TERMINATED | Statistics
28 | pov_ciment_demo | 2007-01-22 14:13:12 | TERMINATED | Statistics
27 | pov_ciment_demo | 2007-01-22 13:53:10 | TERMINATED | Statistics
26 | pov_ciment_demo | 2007-01-22 12:40:13 | TERMINATED | Statistics
25 | pov_ciment_demo | 2007-01-22 12:29:29 | TERMINATED | Statistics
24 | pov_ciment_demo | 2007-01-22 12:10:15 | TERMINATED | Statistics
23 | pov_ciment_demo | 2007-01-21 11:14:23 | TERMINATED | Statistics
22 | pov_ciment_demo | 2007-01-21 10:53:47 | TERMINATED | Statistics

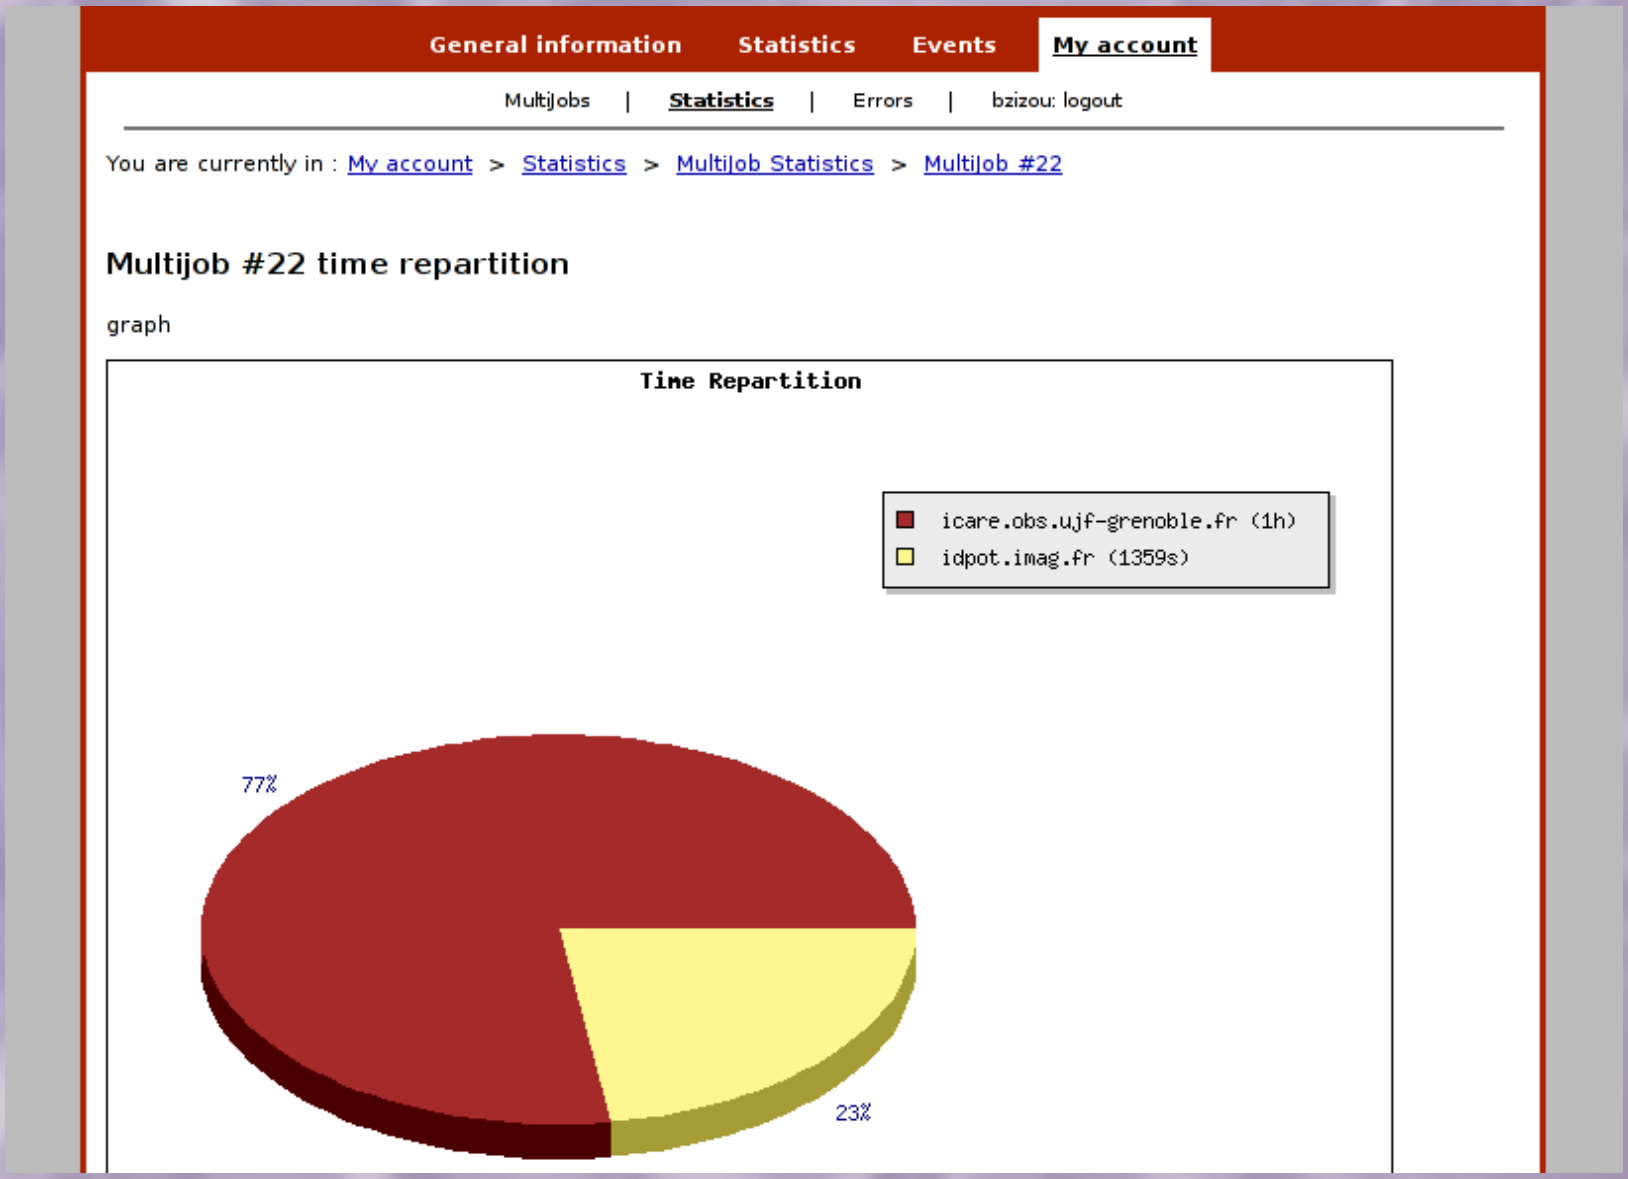
General information
Statistics
Events
My account
MultiJobs | Statistics | Errors | bzizou: logout
You are currently in : My account > Statistics > MultiJob Statistics > MultiJob #22
Multijob #22 time repartition
graph
Time Repartition
icare.obs.ujf-grenoble.fr (1h)
idpot.imag.fr (1359s)
77%
23%

The POVRAY example

- Results collecting:

```
scp "bzizou@clavicule.imag.fr:/home/cigri/results/bzizou/32/*" .

1.idpot.imag.fr.tar.gz
2.idpot.imag.fr.tar.gz
3.idpot.imag.fr.tar.gz
4.icare.obs.ujf-grenoble.fr.tar.gz
5.idpot.imag.fr.tar.gz
6.icare.obs.ujf-grenoble.fr.tar.gz
```

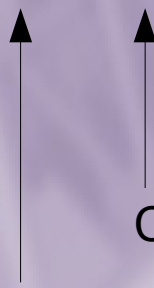

Cluster name

Collect number (1 collect every 30min)

```
# Untar:
for i in *.tar.gz; do tar zxf $i; done
```

The POVRAY example

• Results exploitation:
(we want to put the frames together to make an animation)

```
png2yuv -f 25 -Ip -L0 -j landscape_%04d.png | \
    mencoder -ovc lavc -lavcopts \
    vcodec=mpeg2video:vbitrate=9000 \
    -o landscape_1280x768.m2v -
```

```
mplayer -fs landscape_1280x768.m2v
```