


Application Domain Accounting

J. López Cacheiro, A. Simón, C. Fernández, R. Díez, S. Díaz, A. Rodríguez, A. Gómez (CESGA)

C. Manuali, A. Laganà (University of Perugia)

Application Domain Accounting

- What does it mean?
- What you can get
- What we can do in EGI
- What do you need?

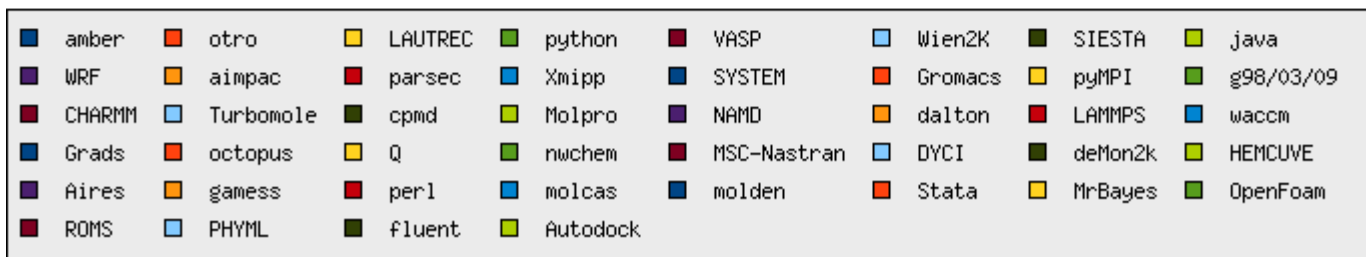





- **Accounting**
 - Job Level
 - Process Level
- **Job level accounting:**
 - Batch system log:
 - SGE: /usr/local/sge/pro/default/common/accounting
 - Grid log:
 - LCG-CE: /opt/edg/var/gatekeeper/grid-jobmap_<date>
 - CreamCE: /opt/glite/var/log/accounting/blahp.log-<date>
- **Application Domain Accounting**
 - Additional information required about the processes run by the jobs




Consumo del día: Miércoles, 07/04/2010

Aplicaciones	Jobs	Tiempo			Memoria	I/O (Chars)	R/W (Blocks)
		sys	usr	elp			
ft							
aimpac	9	9.39	23.49	32.92	17.49	0	0
amber	10625	1,050.50	8,189.99	9,286.60	86.52	0	0
CHARMM	64	6.19	518.99	526.10	26.22	0	0
g98/03/09	59	9.16	1,146.69	449.29	733.28	0	0
Grads	31	0.23	8.73	11.08	8.51	0	0
Gromacs	19	69.07	367.13	439.95	75.82	0	0
LAMMPS	66	1.27	3,176.97	3,182.44	2.57	0	0
molden	1	0.00	0.04	1.13	40.94	0	0
Molpro	68	5.61	434.53	449.15	267.84	0	0
NAMD	12	0.02	0.90	0.94	3.25	0	0
octopus	21	7.87	220.85	229.20	55.22	0	0
otro	1253	618.38	10,986.42	11,620.89	390.08	0	0
parsec	176	2.90	1,022.10	1,027.53	22.08	0	0
python	3	0.15	0.30	119.76	0.84	0	0
SYSTEM	1	0.00	0.04	0.04	33.26	0	0
VASP	92	34.93	419.32	454.82	38.28	0	0
Xmipp	657	20.81	1,929.10	6,784.78	3.09	0	0







Configuración de periodo de tiempo de las tablas.

<input type="checkbox"/>	the process executed fork, but did not exec(AFORK)
<input type="checkbox"/>	the process used super-user privileges(ASU)
<input type="checkbox"/>	the process was killed by a signal(AXSIG)
<input type="checkbox"/>	the process dumped core(ACORE)
<input checked="" type="radio"/> semanal	<input type="radio"/> mensual <input type="radio"/> anual
<input type="button" value="Continuar"/>	

Listado de salidas erróneas							
Aplicaciones	Jobs	Tiempo			Memoria	IO (chars)	RW (blocks)
		sys	usr	elp			
ft							
- proaimv-4.exe	226	156.81	318.83	476.26	17.49	0	0
TOTAL aimpac	226	156.81	318.83	476.26	17.49	0	0
- pmemd	3	25.70	114.62	142.63	12.51	0	0
- ptraj	2	0.00	0.22	0.23	1.54	0	0
- sander.MPI	60486	7,401.99	58,815.70	66,538.59	109.96	0	0
TOTAL amber	60491	7,427.69	58,930.54	66,681.46	109.96	0	0
- charmm	300	43.66	3,666.25	3,715.00	26.22	0	0
TOTAL CHARMM	300	43.66	3,666.25	3,715.00	26.22	0	0
- cpmd.x.mpi	64	10.31	6,376.56	6,399.85	16.71	0	0
TOTAL cpmd	64	10.31	6,376.56	6,399.85	16.71	0	0
- bocimpp.e	1	0.55	151.10	22.59	530.08	0	0
TOTAL DYCI	1	0.55	151.10	22.59	530.08	0	0
- l1002.exe	12	5.27	6,163.33	934.05	801.28	0	0
- l103.exe	5	0.19	58.68	266.22	537.44	0	0
- l1110.exe	3	0.89	466.55	214.50	453.12	0	0



- ExecutingSite
- ExecutingCE
- LocalJobID
- LCGJobID
- LocalUserID
- LCGUserID (UserDN)
- LCGUserVO
(VO/groups/roles)
- StartTime
- StopTime
- ElapsedTime
- BaseCpuTime
- SpecInt2000
- **SpecFloat2000**
- **MemoryReal**
- **MemoryVirtual**



The easy way:

- **VOMS groups can be used to implement it:**
 - Each group corresponds to one application
 - The VO decides the detail level
 - *voms-proxy-init --voms cesga:/cesga/gromacs --order /cesga/gromacs --order /cesga*

```
[jlopez@ui ~]$ voms-proxy-info --all
subject   : /DC=es/DC=irisgrid/O=cesga/CN=javier-lopez/CN=proxy
issuer    : /DC=es/DC=irisgrid/O=cesga/CN=javier-lopez
identity  : /DC=es/DC=irisgrid/O=cesga/CN=javier-lopez
type      : proxy
strength  : 512 bits
path      : /tmp/x509up_u503
timeleft  : 11:59:45
=== VO cesga extension information ===
VO        : cesga
subject   : /DC=es/DC=irisgrid/O=cesga/CN=javier-lopez
issuer    : /DC=es/DC=irisgrid/O=cesga/CN=host/voms.egee.cesga.es
attribute : /cesga/gromacs
attribute : /cesga
timeleft  : 11:59:45
```

The difficult way:

- **Linux Accounting System:**
 - Process accounting log: `/var/account/pacct`
 - It must be activated in each WN
 - All `pacct` files should be collected and merged
 - Difficult to track the UserDN

- **Command name (16 chars)**
- **User ID**
- **Group ID**
- **Process exitcode**
- **Controlling tty**
- **Beginning time**
- **User time**
- **System time**
- **Elapsed time**
- **Chars transferred**
- **Blocks read or written**
- **Average memory usage**
- **Number of swaps**
- **Minor pagefaults**
- **Major pagefaults**

dump-acct /var/account/pacct

command, user time, system time, effective time, uid, gid, memory, time

leiden5 |7172096.0| 843.0|7208960.0|30721|30019|5904.0|Wed Apr 7 14:57:58 2010

[root@compute-0-10 ~]# **sa**

28142 18226.22re 1205.33cp 1036k

1 1201.49re 1195.49cp 1476k leiden5

2 9.50re 9.46cp 363k TinyMC

62 2.32re 0.08cp 1471k globus-url-copy

...

Even more difficult:

- Develop our own sensors

Additionally:

- Provide an API for VOs to publish information
- Store job exit status
- Store memory consumed



- We are working with CompChem VO and CCMST SSC defining the initial requirements
- What other VOs need?
- What NGIs need?
- What EGI needs?

Please send your feedback to:
grid-admin@cesga.es

