Monitoring ATLAS production on LCG

The use of GridICE for job monitoring in ATLAS production

Guido Negri – INFN Cnaf

GridICE

Monitoring tool developed by INFN Grid-it

It provides

- **monitoring over the resources of LCG sites**
- **monitoring of jobs running on sites batch systems**

Why GridICE

While for a single physicist running production the information coming from the BDII can be enough, for coordinating the production an accurate job monitoring is needed, in order to know:

- **where jobs are running**
- **CPU time consumed by jobs**
- **wallclock time consumed by jobs**

plus some few other information

Integration of ATLAS tools

MAGOO is a web tool monitoring ATLAS production by gathering informations coming from

- **ATLAS production (Oracle) database**
- **ATLAS BDII**
- **GridICE (PgSQL database)**

From GridICE, MAGOO extracts informations about jobs running on LCG
The contact point between ATLAS production database and jobs in local resource management systems is the LCG job ID, which is provided by GridICE: through this ID we can get all necessary informations on a job, both on its nature and type (from ATLAS prod DB) and on its current status in the remote batch queue

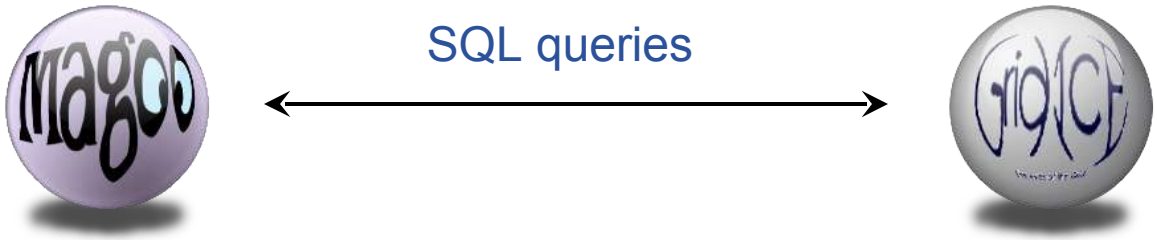

A snapshot from MAGOO

MAGOO gathers information from a GridICE server specifically deployed for the ATLAS VO.

The “Jobs per site” view in Magoo reports the number of running and waiting ATLAS jobs on sites that have GridICE sensor correctly installed and configured.

Clicking on the number of running jobs, you can see the informations provided by GridICE

Dataset overview | Jobs per site | Select jobs | Your jobs | Available sites | Job list match

Site	Running Jobs	Queued Jobs	Successful last 48h	Failed last 48h
grid.sinica.edu.tw	1	0	0	0
bo.infn.it	0	0	0	0
ba.infn.it	0	0	0	0
pg.infn.it	0	0	0	0
lnl.infn.it	34	0	44	0
ific.uv.es	1	0	0	0
roma2.infn.it	0	0	0	0
itep.ru	0	0	0	0
cnb.uam.es	0	0	0	0
pd.infn.it	0	0	0	0
ciemat.es	0	0	0	0

Status	LocalOwner	LocalHost	CPUTime	WallTime	CreationTime	StartTime	RamUsed	VirtualUsed
Running	atlassgm	t2-wn-09	00:00:13	00:06:05	Wed Feb 16 23:36:05 2005	Wed Feb 16 23:36:05 2005	14680	27532
https://lxn1177.cern.ch:9000/eQEEXMWQdTzs94xMw58iZw								

Status	LocalOwner	LocalHost	CPUTime	WallTime	CreationTime	StartTime	RamUsed	VirtualUsed
Running	atlassgm	atlfarm012	05:13:16	05:35:15	Wed Feb 16 20:44:20 2005	Wed Feb 16 20:44:20 2005	317316	604908
https://lcg2rb2.ific.uv.es:9000/n2orv7mDXtWTIYfC9xj6Qg								

Status	LocalOwner	LocalHost	CPUTime	WallTime	CreationTime	StartTime	RamUsed	VirtualUsed
Running	atlassgm	grid021	03:58:15	05:35:54	Wed Feb 16 20:44:19 2005	Wed Feb 16 20:44:19 2005	320396	603884
https://lcg2rb2.ific.uv.es:9000/Xpq0j4OWr038QLax4xstcQ								

Status	LocalOwner	LocalHost	CPUTime	WallTime	CreationTime	StartTime	RamUsed	VirtualUsed
Running	atlassgm	grid021	01:45:25	05:35:55	Wed Feb 16 20:44:19 2005	Wed Feb 16 20:44:19 2005	336024	596032
https://lcg2rb2.ific.uv.es:9000/eR4FYfTxU4creSEs05TMTQ								

Clicking on the job ID you can get an overview of the job from the prod DB

What we still lack

Unfortunately, we still cannot rely entirely on MAGOO and GridICE for job monitoring because the GridICE sensor (the one responsible of collecting and publishing data on running jobs) is NOT YET installed on many production sites.

GridICE team has developed some easy installation and configuration tools for this sensor, also thanks to the cooperation by some LCG sitemanagers who tried out the installation and provided GridICE developers with useful feedback.

Hopefully this will encourage a wider deployment of the sensor.