



Departament de Llenguatges i Sistemes Informàtics

S-GAE

SunGrid Graphical Accounting Engine

<http://rdlab.lsi.upc.edu>
rdlab@lsi.upc.edu

Gabriel Verdejo Àlvarez
Fernando Galindo Pascual

October 2011

Contents

1. Who we are
2. The RDlab cluster
3. Accounting
4. Existing software
5. The S-GAE application
6. S-GAE quick view
7. How S-GAE is released
8. Live Demo

1. Who we are

- RDIlab: Research and Development Laboratory.
- Belongs to LSI department at UPC University.
- Created on fall 2010.
- Currently composed by 5 people.

1. Who we are II: Our numbers

- Over 130 computer servers
 - > 70 nodes RDIlab HPC cluster
 - > 60 project servers
- TIC support for 8 research groups at LSI

ALBCOM

KEMLG

GRPLN

GIE

LOGPROG

LARCA

MOVING

SOCO

- CEE projects & Enterprise Agreements

1. Who we are III: Our homepage

The screenshot shows a Mozilla Firefox browser window with the title "Benvinguts a RDLab - Mozilla Firefox". The address bar displays "http://rdlab.lsi.upc.edu/". The page content includes a large blue header with the "RDLab" logo and the text "Departament de Llenguatges i Sistemes Informàtics". Below the header is a navigation menu with links: "INICI", "QUI SOM?", "SERVEIS", "CONTACTEU AMB NOSALTRES", and "MAPA WEB". A main content area features a large photograph of the "Edifici Omega" at the "Campus Nord de la UPC". To the right of the photo is a "Google Search" box with a search bar and a "Search" button. Below the search box is an "Avisos" section with a red link for "HPC Sys Admin Meeting '11" and the text "Antics". At the bottom of the page, there are two RSS feed sections: "RSS De La Sala De Premsa UPC" and "RSS De La Xarxa TERENA". The "RSS De La Sala De Premsa UPC" section lists three items: "Josep Díaz, Premi Nacional d'Informàtica 2011", "Agbar, la UPM i la UPC impulsen projectes de formació", and "CA Technologies obre un nou centre de recerca europeu al Campus Nord de la UPC". The "RSS De La Xarxa TERENA" section lists three items: "Hacking Away at the Internet's Security", "GLIF community finds rhythm in samba capital", and "Congress of trauma surgery and forensic medicine".

Benvinguts a RDLab - Mozilla Firefox

Archivo Editar Ver Historial Marcadores Herramientas Ayuda

Benvinguts a RDLab

http://rdlab.lsi.upc.edu/

RDLab

Departament de Llenguatges i Sistemes Informàtics

INICI QUI SOM? SERVEIS CONTACTEU AMB NOSALTRES MAPA WEB

|A+ a-

Edifici Omega
Campus Nord de la UPC

Google Search

Google Custom Search

Web rdlab.lsi.upc.edu

Search

Avisos

HPC Sys Admin Meeting '11

Antics

RSS De La Sala De Premsa UPC

- Josep Díaz, Premi Nacional d'Informàtica 2011
- Agbar, la UPM i la UPC impulsen projectes de formació
- CA Technologies obre un nou centre de recerca europeu al Campus Nord de la UPC

RSS De La Xarxa TERENA

- Hacking Away at the Internet's Security
- GLIF community finds rhythm in samba capital
- Congress of trauma surgery and forensic medicine

2. The RDIlab Cluster: Our size

- More than 70 nodes.
- More than 560 execution cores.
- More than 1 Terabytes of RAM.
- More than 10 Terabytes of disk space.
- More than 120 users.
- More than 339.000 executed jobs since January 1st!

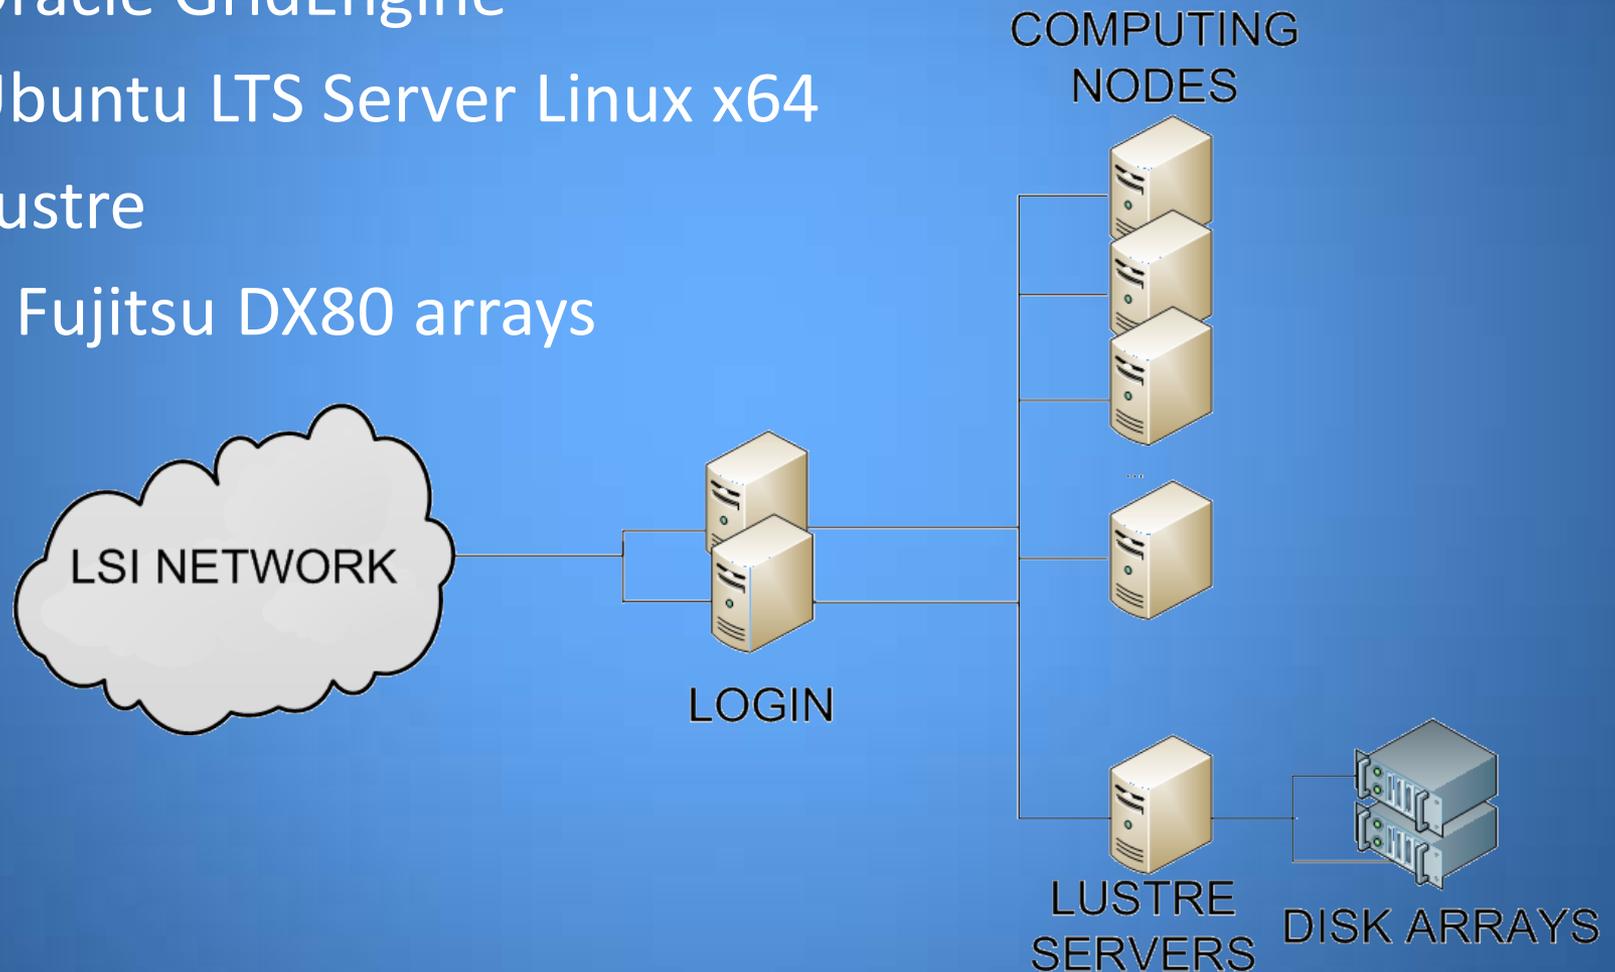
2. The RDIlab Cluster II: Facilities

- Located at UPC's Data Processing Centre (CPD).
- More than 250m² surface area.
- Back-up generator with complete autonomy.



2. The RDLab Cluster III: Software

- Oracle GridEngine
- Ubuntu LTS Server Linux x64
- Lustre
- 2 Fujitsu DX80 arrays



3. Accounting

- User activity

Large number of users

+ A load of jobs per user

LOTS of raw accounting data

- We needed to monitor all this information
- Oracle GridEngine saves this data in a plain text file
- We needed an interface to get easy access to it

3. Accounting II

- Objectives:
 - Monitor and control accounting information
 - Avoiding *Command Line Interface*
 - Form-based customization filters
 - Show aggregated values (avg., %,...)
 - Customized view for cluster users, queue managers and cluster administrators.

4. Existing software

- ARCo:
 - Excessively large database size.
 - SQL based interface for queries.
- Gridsafe: <http://gridsafe.sourceforge.net/index.html>
 - Large list of requirements
 - Poor documentation
 - Not specifically designed for SGE
- PHPQstat: <http://sourceforge.net/projects/phpqstat/>
 - Only a qstat interface.

4. Existing software II

- None of them suited our expectations
- We decided to develop our own product

SunGrid Graphical Accounting Engine

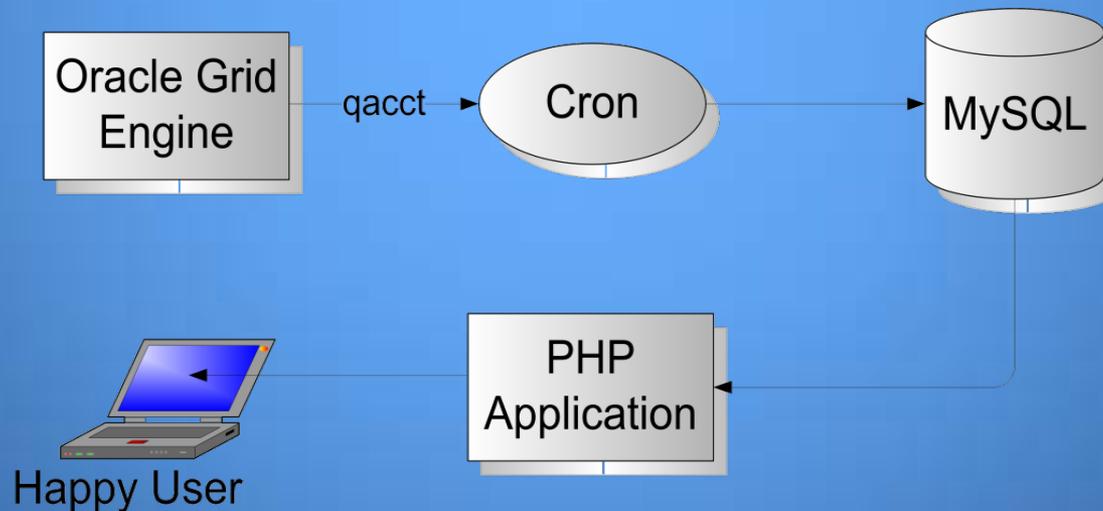
- Trivia:
 - Oracle Engine is not SunGrid anymore.
 - Why “*s-gae*” and not simply “*sgae*”...?

6. S-GAE quick view I: How it works

- Parses Oracle GridEngine accounting data.
 - Periodically (in batch mode): cron
 - On demand: administrator interface
- Compacts and processes data.
- Store it into a MySQL database
 - Group data according to queue name
 - Group by year / month of submission

6. S-GAE quick view II: How it works

- PHP gets data from the database
 - Ready for graphs generation (no further processing).
- Customize results through filters and show graphs

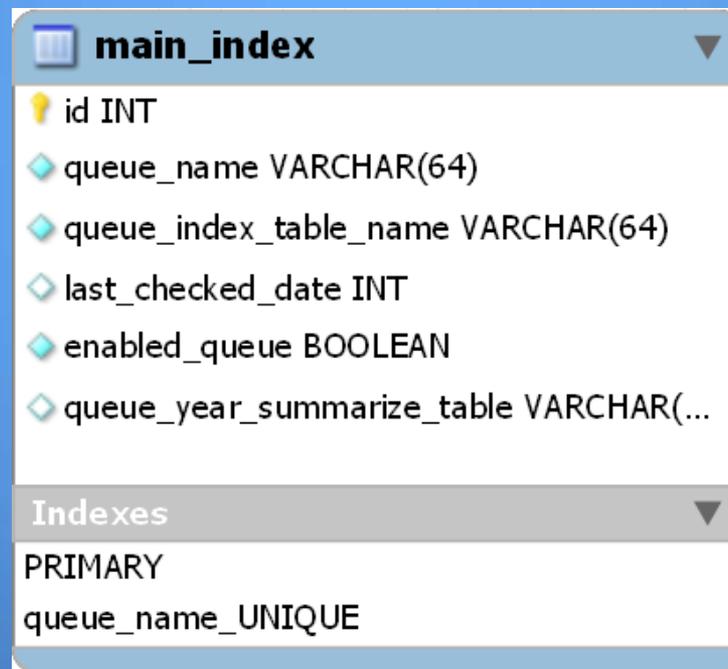


6. S-GAE quick view III: Database

- Most of work is done by the Database Engine
- Data processing is done by Procedures
 - Creates a view with the data needed
 - Calculates averages, top ten orders, etc.
 - Leave result in a single-row table
- PHP calls stored procedures and gets the result
- This reduces data transmission and PHP workload

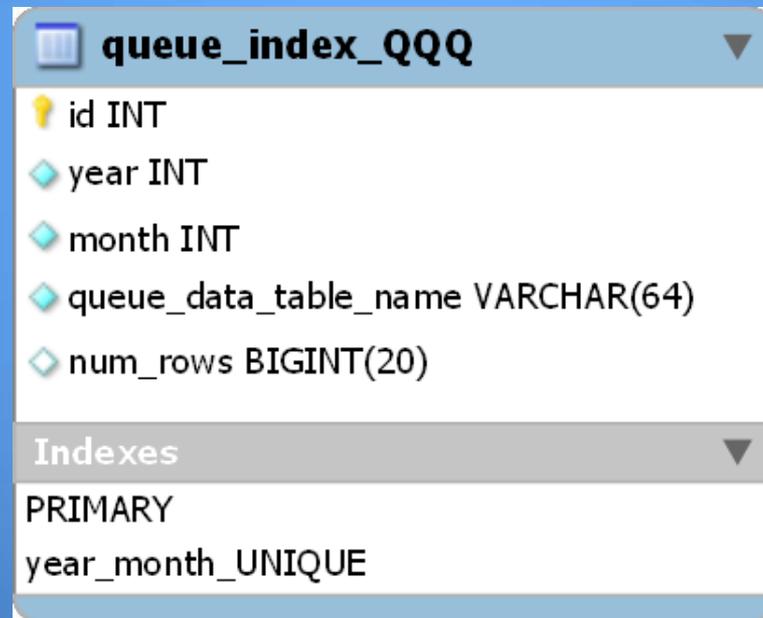
6. S-GAE quick view IV: Database

- Main index: 1 table
 - Name of queues, enabled, last check date



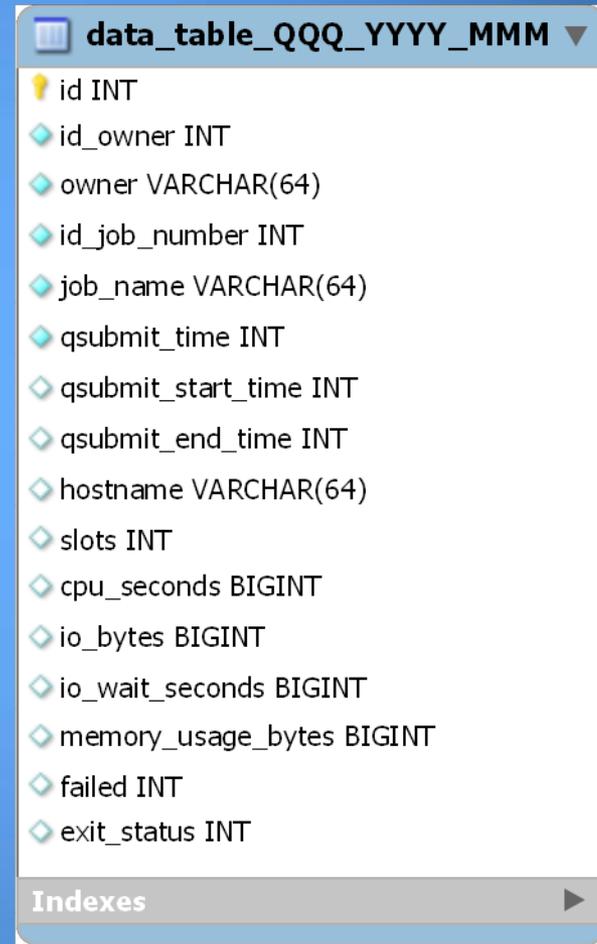
6. S-GAE quick view V: Database

- Queue index: 1 table per queue
 - Name of data tables by year and month



6. S-GAE quick view VI: Database

- Data table:
 - 1+ per queue
 - Parsed data



The screenshot shows a database table schema for 'data_table_QQQ_YYYY_MMM'. The table has the following columns:

Column Name	Column Type
id	INT
id_owner	INT
owner	VARCHAR(64)
id_job_number	INT
job_name	VARCHAR(64)
qsubmit_time	INT
qsubmit_start_time	INT
qsubmit_end_time	INT
hostname	VARCHAR(64)
slots	INT
cpu_seconds	BIGINT
io_bytes	BIGINT
io_wait_seconds	BIGINT
memory_usage_bytes	BIGINT
failed	INT
exit_status	INT

Below the column list, there is a section for 'Indexes' with a right-pointing arrow.

6. S-GAE quick view VII: Database

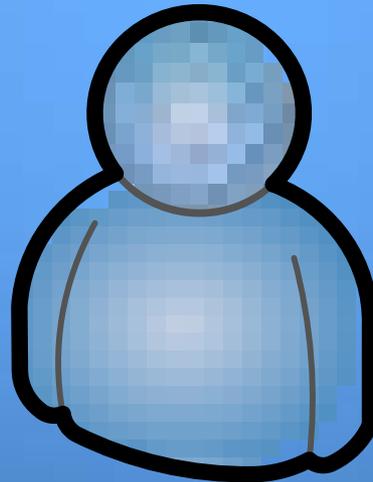
- Year summarize charts: 1 per queue

```
queue_year_summarize_charts
id INT
year INT
last_update_time INT
jobs_execution_time_longest INT
jobs_execution_time_average INT
jobs_execution_time_shortest INT
jobs_execution_time_1d INT
jobs_execution_time_7d INT
jobs_execution_time_31d INT
jobs_execution_time_xxd INT
jobs_success INT
jobs_failure INT
jobs_memory_maximum BIGINT
jobs_memory_average BIGINT
jobs_memory_minimum BIGINT
jobs_memory_xxgb INT
jobs_memory_8gb INT
jobs_memory_4gb INT
jobs_memory_1gb INT
```

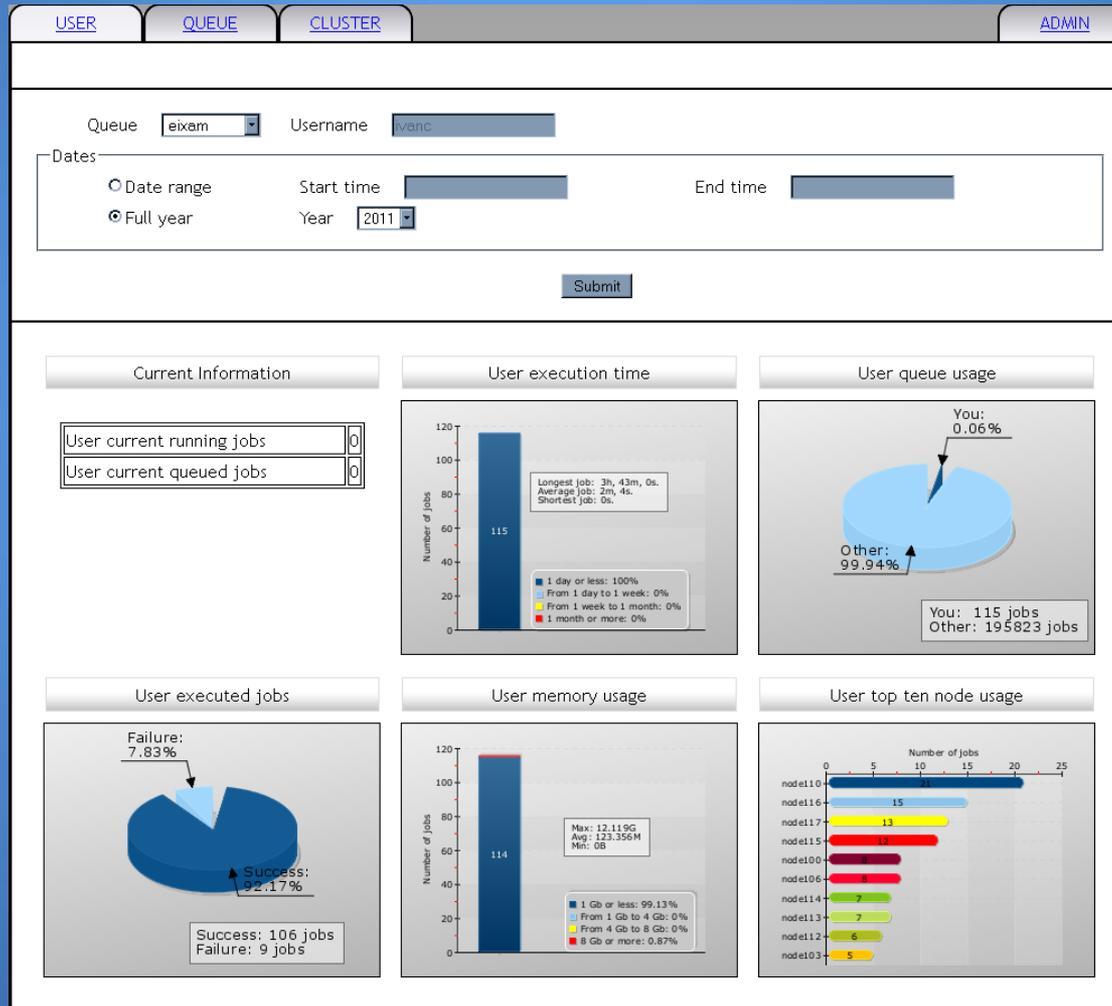
```
hostname_usage_top1 INT
..... INT
hostname_usage_top10 INT
hostname_usage_top1_host VARCHAR(64)
..... VARCHAR(64)
hostname_usage_top10_host VARCHAR(64)
jobs_sent_user_top1 INT
..... INT
jobs_sent_user_top10 INT
jobs_sent_user_top1_owner VARCHAR(64)
..... VARCHAR(64)
jobs_sent_user_top10_owner VARCHAR(64)
jobs_memory_usage_top1 BIGINT
..... INT
jobs_memory_usage_top10 BIGINT
jobs_memory_usage_top1_owner VARCHAR(64)
..... VARCHAR(64)
jobs_memory_usage_top10_owner VARCHAR(...)
jobs_execution_time_top1 BIGINT
..... INT
jobs_execution_time_top10 INT
jobs_execution_time_top1_owner VARCHAR(64)
..... VARCHAR(64)
jobs_execution_time_top10_owner VARCHAR(...)
```

6. S-GAE quick view VIII: Roles

- Common user
- Queue administrator
- Cluster and Application administrator



6. S-GAE quick view IX: Interface



7. How S-GAE is released

- GNU General Public License v3
- Visit our homepage <http://rdlab.lsi.upc.edu/s-gae>

The screenshot shows the SunGrid Graphical Accounting Engine (S-GAE) homepage. The browser window title is "SunGrid Graphical Accounting Engine - Mozilla Firefox". The address bar shows the URL <http://rdlab.lsi.upc.edu/index.php/ca/serveis/s-gae.html>. The page header features the RDLab logo and the text "Departament de Llenguatges i Sistemes Informàtics". A navigation menu includes "INICI", "QUE SOM?", "SERVEIS", "CONTACTEU AMB NOSALTRES", and "MAPA WEB". The main content area is titled "S-GAE" and "SunGrid Graphical Accounting Engine by RDLab". Below this, a section titled "What is s-gae?" explains that S-GAE is a web application designed to display accounting information generated by Oracle Grid Engine (formerly SunGrid Engine). It mentions that the data is stored in a database and can be displayed as eye-candy charts grouped by user, queue, or full cluster. To the right, there is a "Google Search" box and an "Avisos" section with a link to "HPC SysAdmin Meeting T1". At the bottom of the page, there are several charts and a table displaying performance metrics.

Live demo

CONNECT

S-GAE

- Web browser interface with eye-candy charts
- Form-based data filtering and aggregated values
- User, queue and full cluster statistics
- Licensed under GPL v3

<http://rdlab.lsi.upc.edu/s-gae>