

Sun Grid Engine Package for OSCAR

A Google SoC 2005 Project

Babu Sundaram, Barbara Chapman
University of Houston

Bernard Li, Mark Mayo, Asim Siddiqui, Steven Jones
Canada's Michael Smith Genome Sciences Centre



Department of Computer Science
The University of Houston

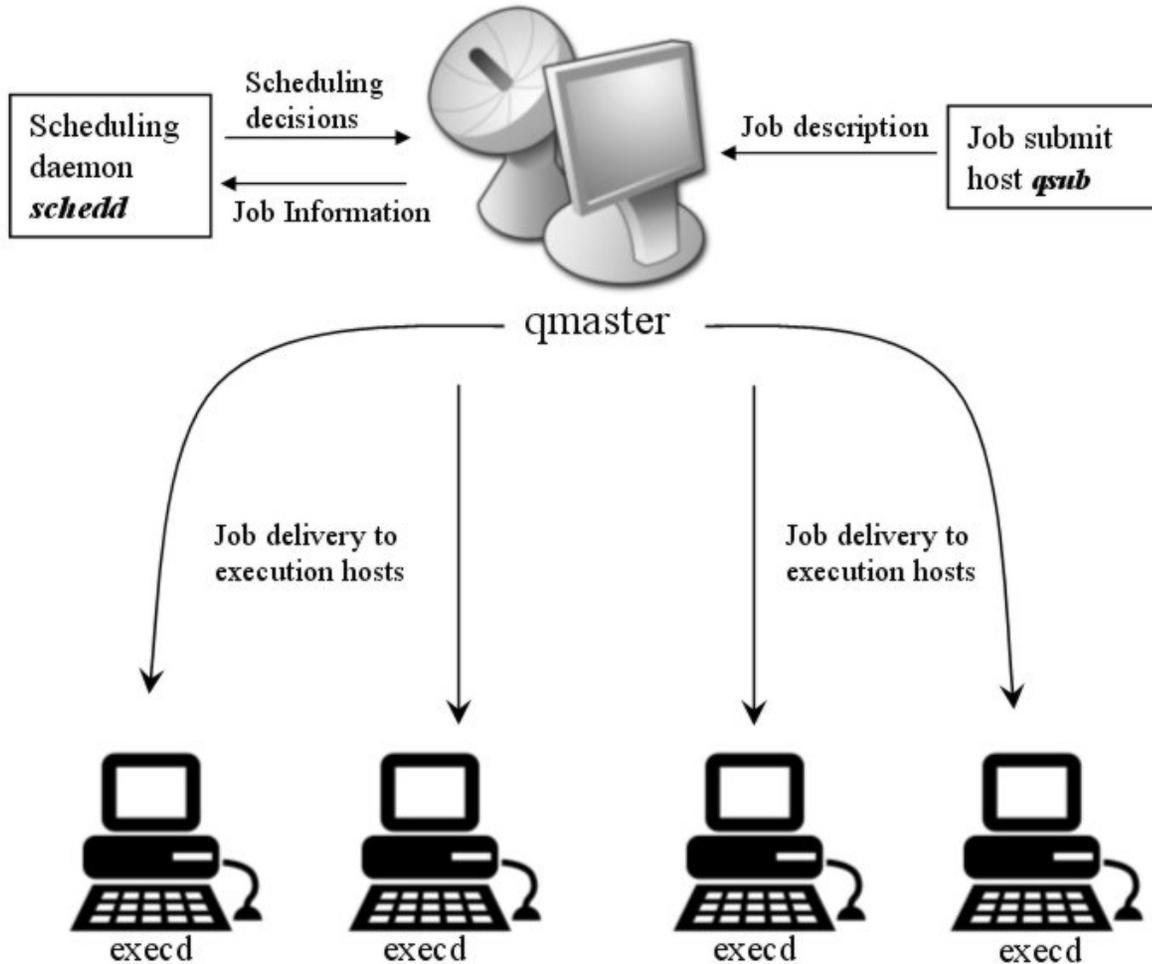
Sun Grid Engine

- **Distributed resource management and batch job queuing software**
- **Increase cluster utilization to maximum**
- **Precise control over resource usage, supports sophisticated scheduling policies**
- **Widely deployed at major institutions**
 - UH (COE) has a SGE cluster (~250 nodes)
- **Open source software, community effort**
 - `gridengine.sunsource.net`



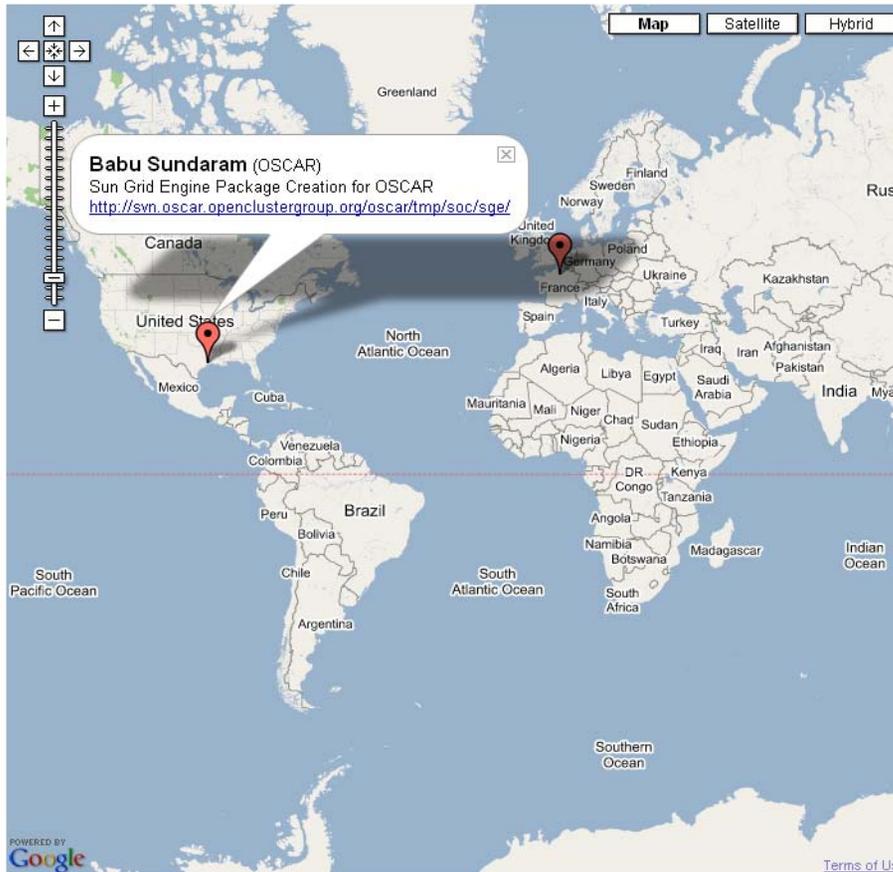
Department of Computer Science
The University of Houston

Typical SGE setup



Summer of Code: Where are they now?

To show you the global reach of the Summer of Code, we've put together a map showing the locations of both our students and mentors around the world (or at least the students and mentors that told us where they are!). Click the checkboxes to see which project participants are where, and click any marker to find out who it belongs to.



Show: All Mentors Students

- Toggle all
- Apache
- Blender
- Codehaus
- Drupal
- Fedora
- FreeBSD
- Gaim
- Gallery
- Google
- Handhelds
- Horde
- Inkscape
- Internet2
- JXTA
- Jabber
- Joomla
- KDE
- LispNYC
- LiveJournal
- Monotone
- Mozdev
- NMap
- NetBSD
- OSCAR
- OpenOffice
- Other
- PSU
- Project Looking Glass
- Python
- Samba
- Semedia
- GNOME
- Mono
- Perl
- Subversion
- Wine
- Ubuntu
- WinLibre
- XWiki

Map Students to Mentors
(Not all lines will be shown since not all students and mentors are mapped)



The OSCAR Project

- *“...a snapshot of the best known methods for building, programming, and using HPC clusters”*
- Easy to install software bundle
- Everything needed to install, build, maintain and use a Linux cluster
- Supports various distros such as Red Hat Enterprise Linux (and clones), Fedora Core, Mandriva Linux on x86, ia64, x86_64 architectures
- <http://oscar.openclustergroup.org>



Department of Computer Science
The University of Houston

What is an OSCAR Package?

```
<packages dir>
```

*** - mandatory**

```
<package_name>
```

```
config.xml* doc RPMS SRPMS scripts testing
```



Department of Computer Science
The University of Houston

OSCAR Package details

- `config.xml` – XML file indicating package details, its version, dependencies (e.g., `sge`, `ksh`) and OS-, client-specific `rpmlists`
- `doc` – Mostly help and `README` files
- `RPMS` – pre-compiled binaries as `RPMs`
- `SRPMS` – to allow building on other platforms
- `testing` – tests after package installation



Department of Computer Science
The University of Houston

OSCAR Package scripts

- **OSCAR framework recognizes a standard set of scripts and they have definitive purpose**

Seq#	Script Name	Description
1	<code>setup</code>	Perform any package setup
2	<code>pre_configure</code>	Prepare package config (dynamic user input)
3	<code>post_configure</code>	Process results from package config
4	<code>post_server_rpm_install</code>	Perform “out of RPM” operations on server
5	<code>post_client_rpm_install</code>	Perform “out of RPM” operations on client
6	<code>post_clients</code>	For configurations with knowledge about nodes
7	<code>post_install</code>	For final config with fully install/booted nodes



Department of Computer Science
The University of Houston

OSCAR Package Configuration

- **configurator.html**
 - page with configuration settings to be used during the "Configure Selected OSCAR Packages" step
- **Values stored in .configurator.values and used by scripts for setup**



Department of Computer Science
The University of Houston

SGE Package for OSCAR

- Lots of interest for SGE OSCAR Package
- Provides an alternative Resource Manager to TORQUE
- Sets up SGE as part of cluster deployment or add-on after initial deployment



Department of Computer Science
The University of Houston

Tasks in SGE package creation

- **Source RPM generation**
- **Binary RPM generation**
 - `Server-`, `client-` and `GUI-specific` RPMs
- **Develop OSCAR configuration and scripts**
- **Implementation, Licensing, Documentation**



Department of Computer Science
The University of Houston

RPM generation for SGE

- **Source RPM generation was our first step**
- **SGE source rpm for version 6.0 update 4**
 - At that time, ScalableSystems had a release ready
 - Now, we have SRPM and RPM based on update 8
- **Some patches were identified earlier on and some were added later for correct compilation**
 - qtssh, inst_sge, aimk, distinst, qmon icons
- **Spec file modification and SGE binary RPM generation**



Department of Computer Science
The University of Houston

Scripts for SGE-OSCAR

- Automates SGE install on the OSCAR cluster
- All perl scripts

- **post_server_install**

- Configures the overall SGE setup; Sets up SGE master with various values for the options
 - SGE_ROOT, CELLNAME, FULLSERVER, GIDRANGE, SPOOLTYPE, PORTS...
 - oscar_cluster.conf is a file that gets generated at this stage to drive "inst_sge -m -auto"
- User input/customization happens at this stage (configurator.html)
- At the end of this step, the qmaster is up and running on the OSCAR head node

- **post_clients**

- Gets executed after clients are defined (not installed)
- Adds clients as admin hosts so they can be setup as exec hosts later
- `get_machine_listing(); then, qconf -ah $hostname;`



SGE OSCAR scripts – cont...

• `post_install`

- All actions that can be done only after a full cluster install happen in this step
- `qmaster` already knows about the clients (from the definition step) and they are already admin hosts
- All settings (`dir: cell_name`) gets tarred and ready to get pushed to the clients during `post_install`
- Cannot assume NFS; So, the `cell_name_dir.tar` gets pushed to the clients and untarred
- Clients now know about the `qmaster` details
- Automated install of `inst_sge -x` (patched in spec); Executed via `cexec` over `ssh`

• `post_server_rpm_uninstall`, `post_client_rpm_uninstall`

- Not much SGE-specific functionality, but there to allow clean SGE uninstall



Department of Computer Science
The University of Houston

Implementation details

- **OSCAR's Subversion repository for code revision control**
 - `http://svn.oscar.openclustergroup.org/oscar/tmp/soc/sge`
- **Initial implementation was on FC2 x86**
- **Basic tools involved: rpm, make, perl, diff/patch**
- **OSCAR-specific code is under GPL; SGE under SISSL**



Department of Computer Science
The University of Houston

Where is the code now?

- Code integrated into **OSCAR** trunk, to be released in 5.0
- Supported by all distributions on x86 and x86_64 (except for Mandriva)
- Parallel Environment integration: LAM/MPI, PVM, MPICH, Open MPI (only setup if parallel libraries are installed)



Department of Computer Science
The University of Houston

Acknowledgements

- Google Inc.,
- OSCAR developers
- SGE developers (Ron, Fritz, Andreas...)
- Chandler Wilkerson, LAN admin, CS, UH
- ScalableSystems



Department of Computer Science
The University of Houston