



Center for Mathematical Modeling  
University of Chile



# HPC 101

Scientific Computing on HPC systems

By

**Juan Carlos Maureira B.**

*<jcm@dim.uchile.cl>*


*Second La Serena School for Data Science: Applied Tools for Astronomy.*  
La Serena – Chile – 17/08/2014

# Overview


- Concepts & Definitions.
- Working with a HPC system.
- Best Practices.
- Hands-on: My first sextraction
- Closing words: The take aways.

# Scientific Computing

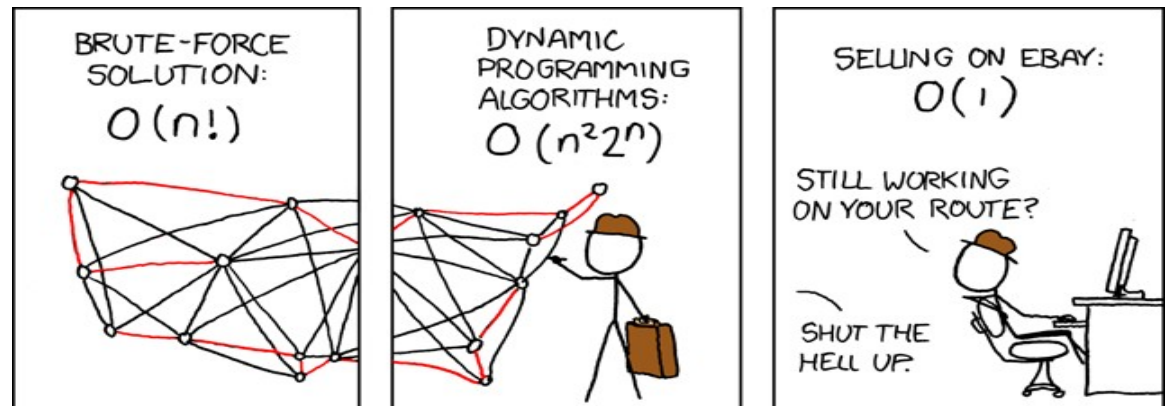
- Simulations



- Data Analysis



- Computational Optimization



# Concepts & Definitions

# HPC system Architecture

- Areas

- Computing, Storage, Support, Networking

- Servers roles

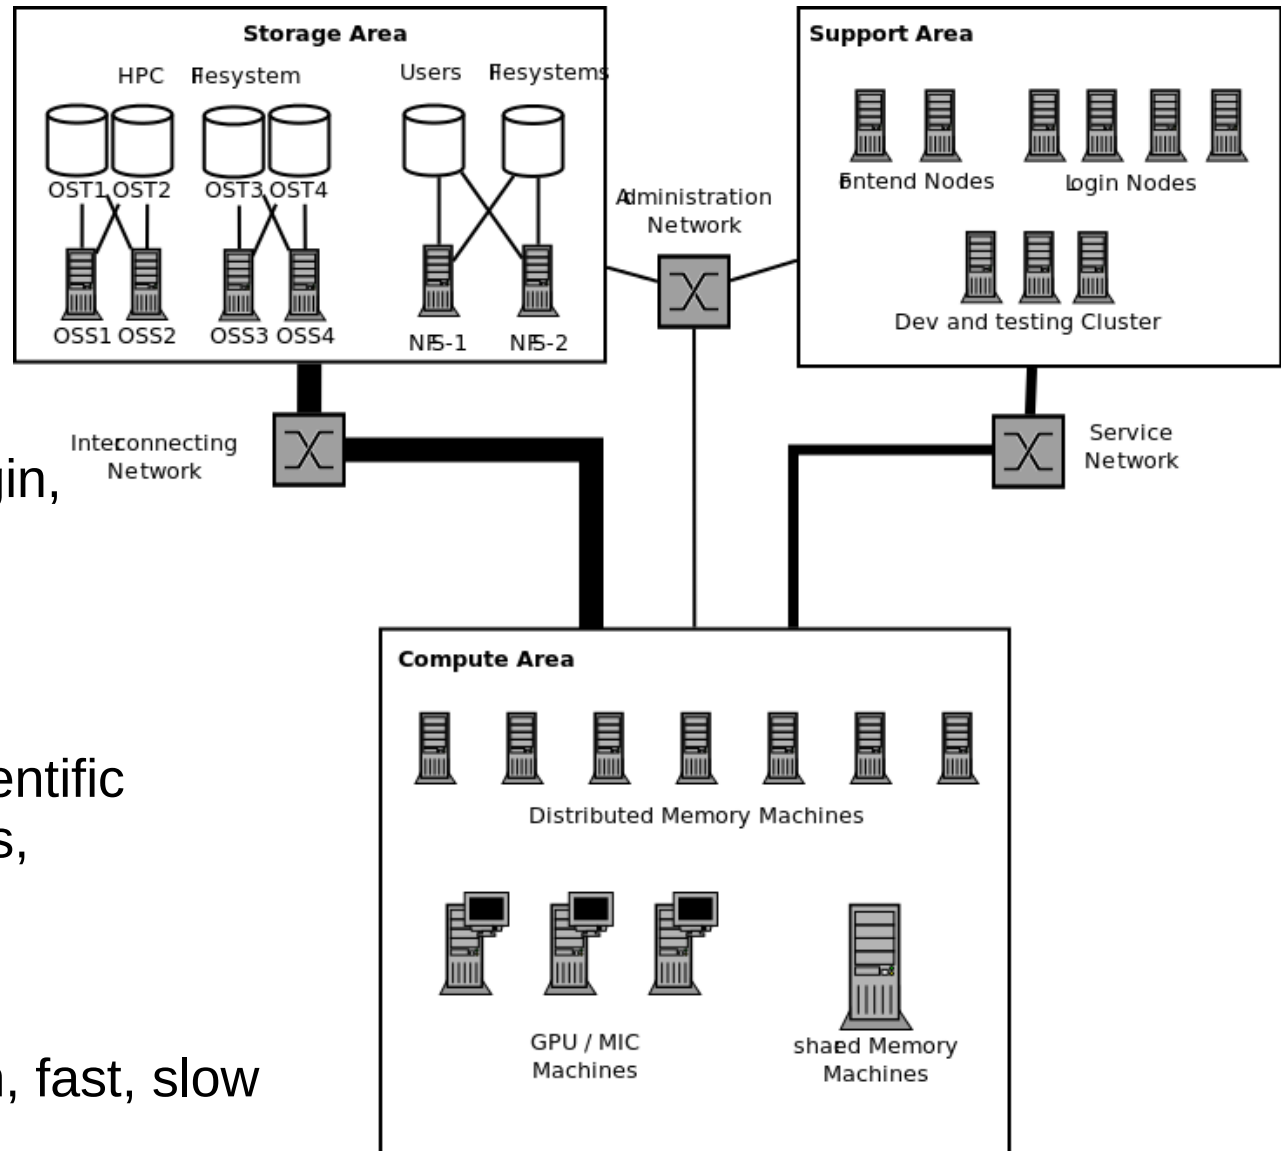
- Compute, frontend, login, storage, backup, devel monitoring, etc.

- Software


- Operating System, scientific software, analysis tools, libraries, etc.

- Storage


- Local, working, scratch, fast, slow




# Distributed and Shared Memory Systems



Shared Memory




Distributed Memory



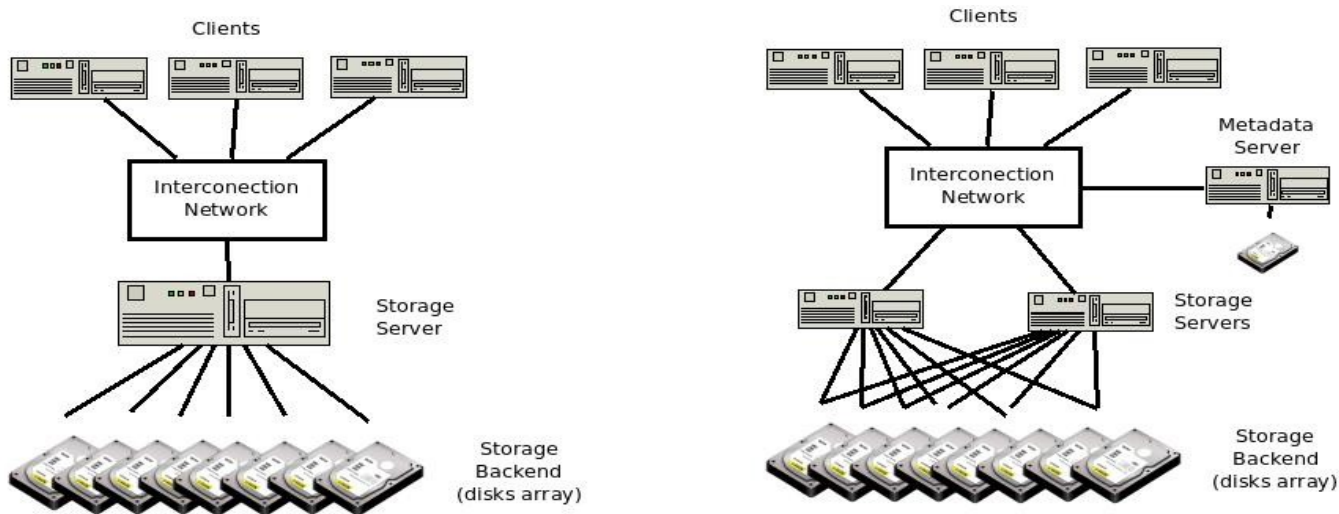
Distributed Shared Memory

# Interconnects


- Ethernet
  - latency ~ 0.05 ms
  - Throughput ~ 10 Gbps
- Infiniband
  - latency ~5 usec
  - Throughput ~ 40/56 Gbps
- QPI / NUMA
  - Latency ~ 100 nsec
  - Throughput ~ 100 - 200 Gbps




# File-systems Types



Serial



Distributed



Parallel

- Serial
  - NFS, ZFS
- Distributed
  - pNFS
  - GFS
  - Gluster
- Parallel
  - Lustre
  - GPFS



# Storage Layouts


- **Working (\$Home)**

- Safe and **Slow** storage.
- Cheep
- Bad for I/O



- **Scratch**

- Unsafe and **Fast** storage
- Expensive
- Volatile and great for I/O

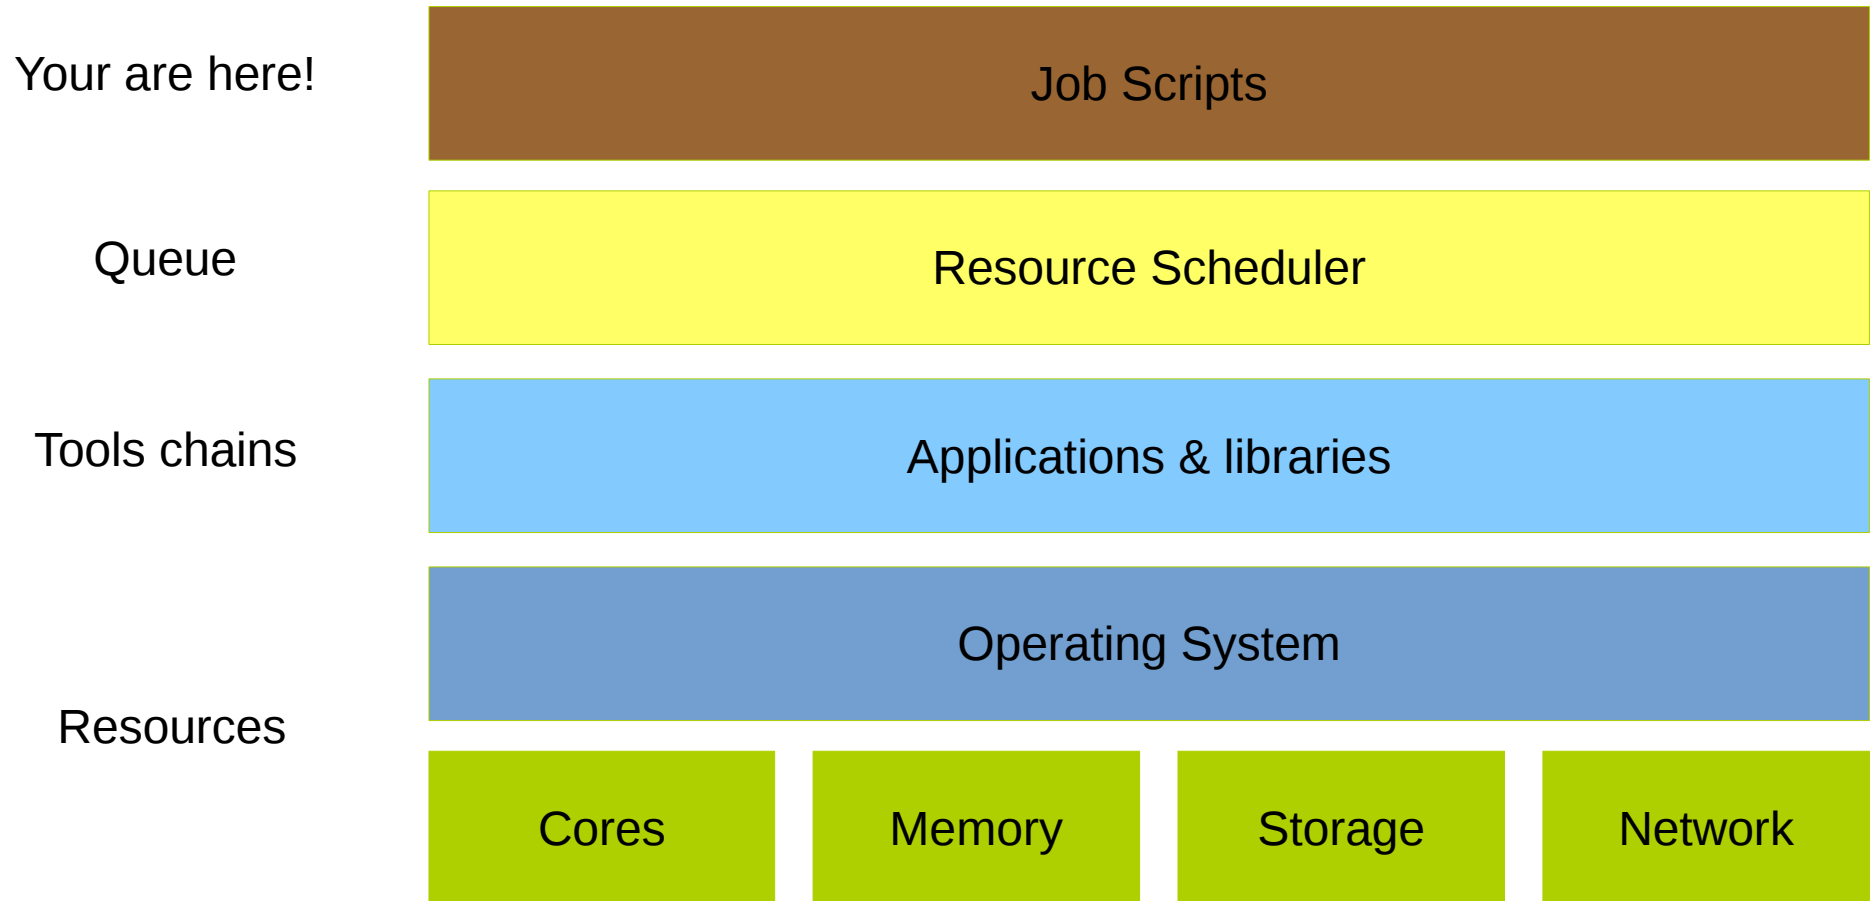


- **Archiving**

- disaster-proof storage
- **Incredible slow** (random) access
- Backup Policies



# Software Layout



# Tool Chains

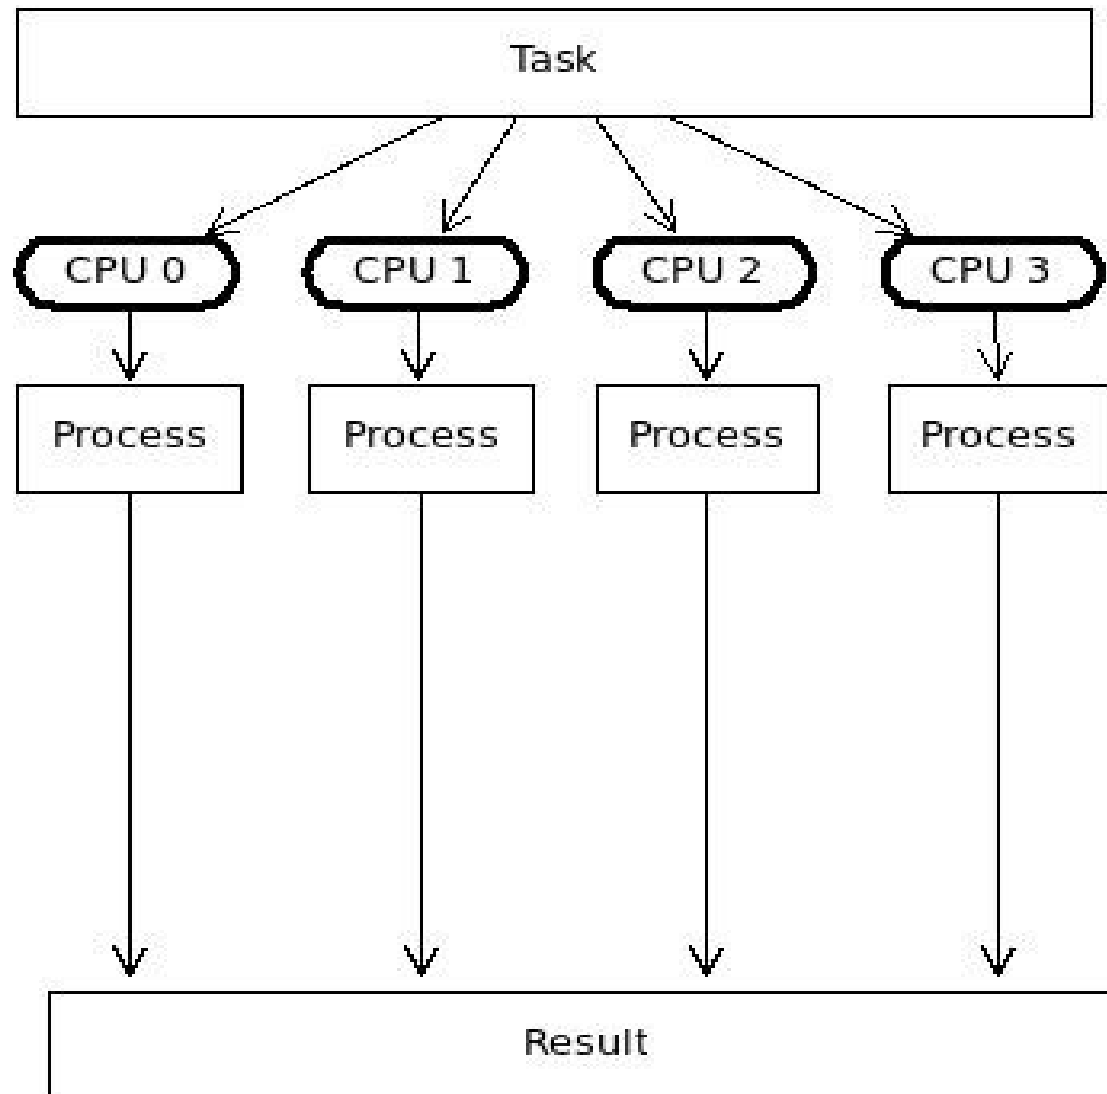
- Set of self-standing libraries and applications to perform a class of jobs. (e.g. astro, bioinfo, optimization, etc).
- System wide (one for all).
  - Compiled and Installed by admins.
- User Space (each one has its own).
  - Compiled and installed by the user in their homes.

# Resource Schedulers

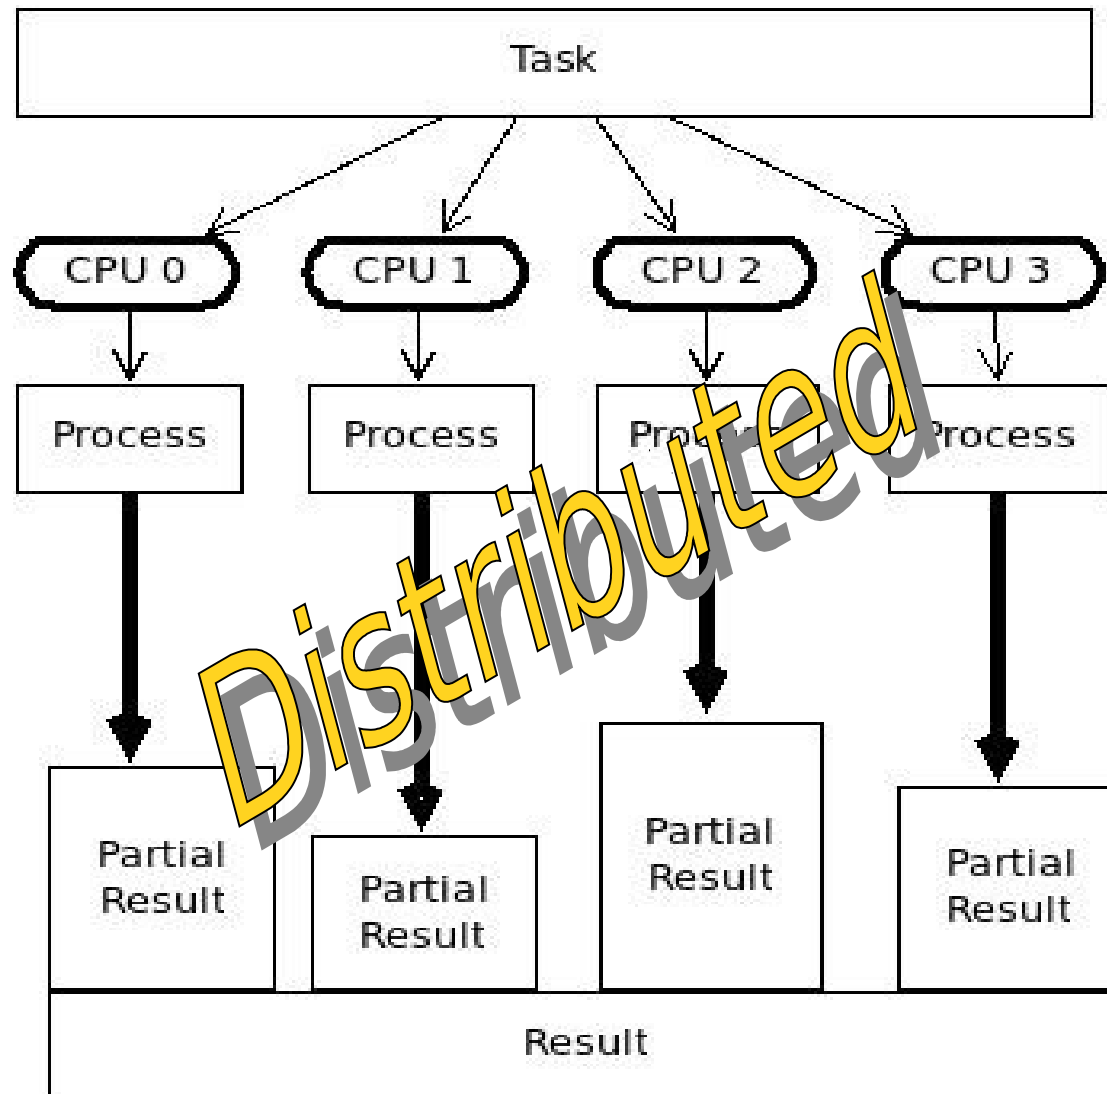
- **Scheduler:** allocate resources to perform a job.
- **Job:** set of instructions and resources to perform a task.
- **Task:** involves preparing the environment and input data needed to run an application.

**Resource specifications**  
**+**  
**Instructions to perform a task**

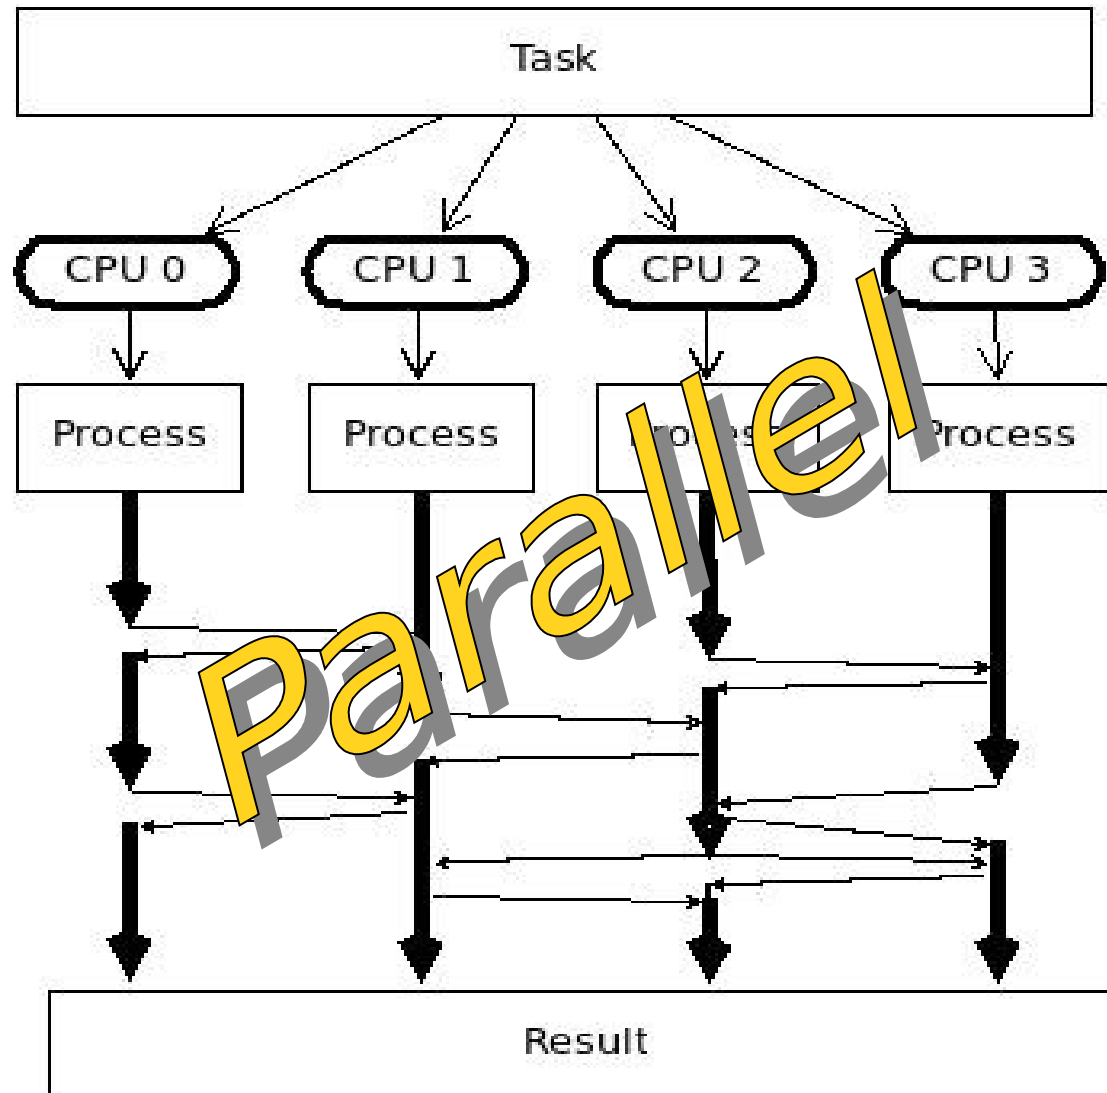
# Jobs: Parallel v/s Distributed



# Jobs: Parallel v/s Distributed



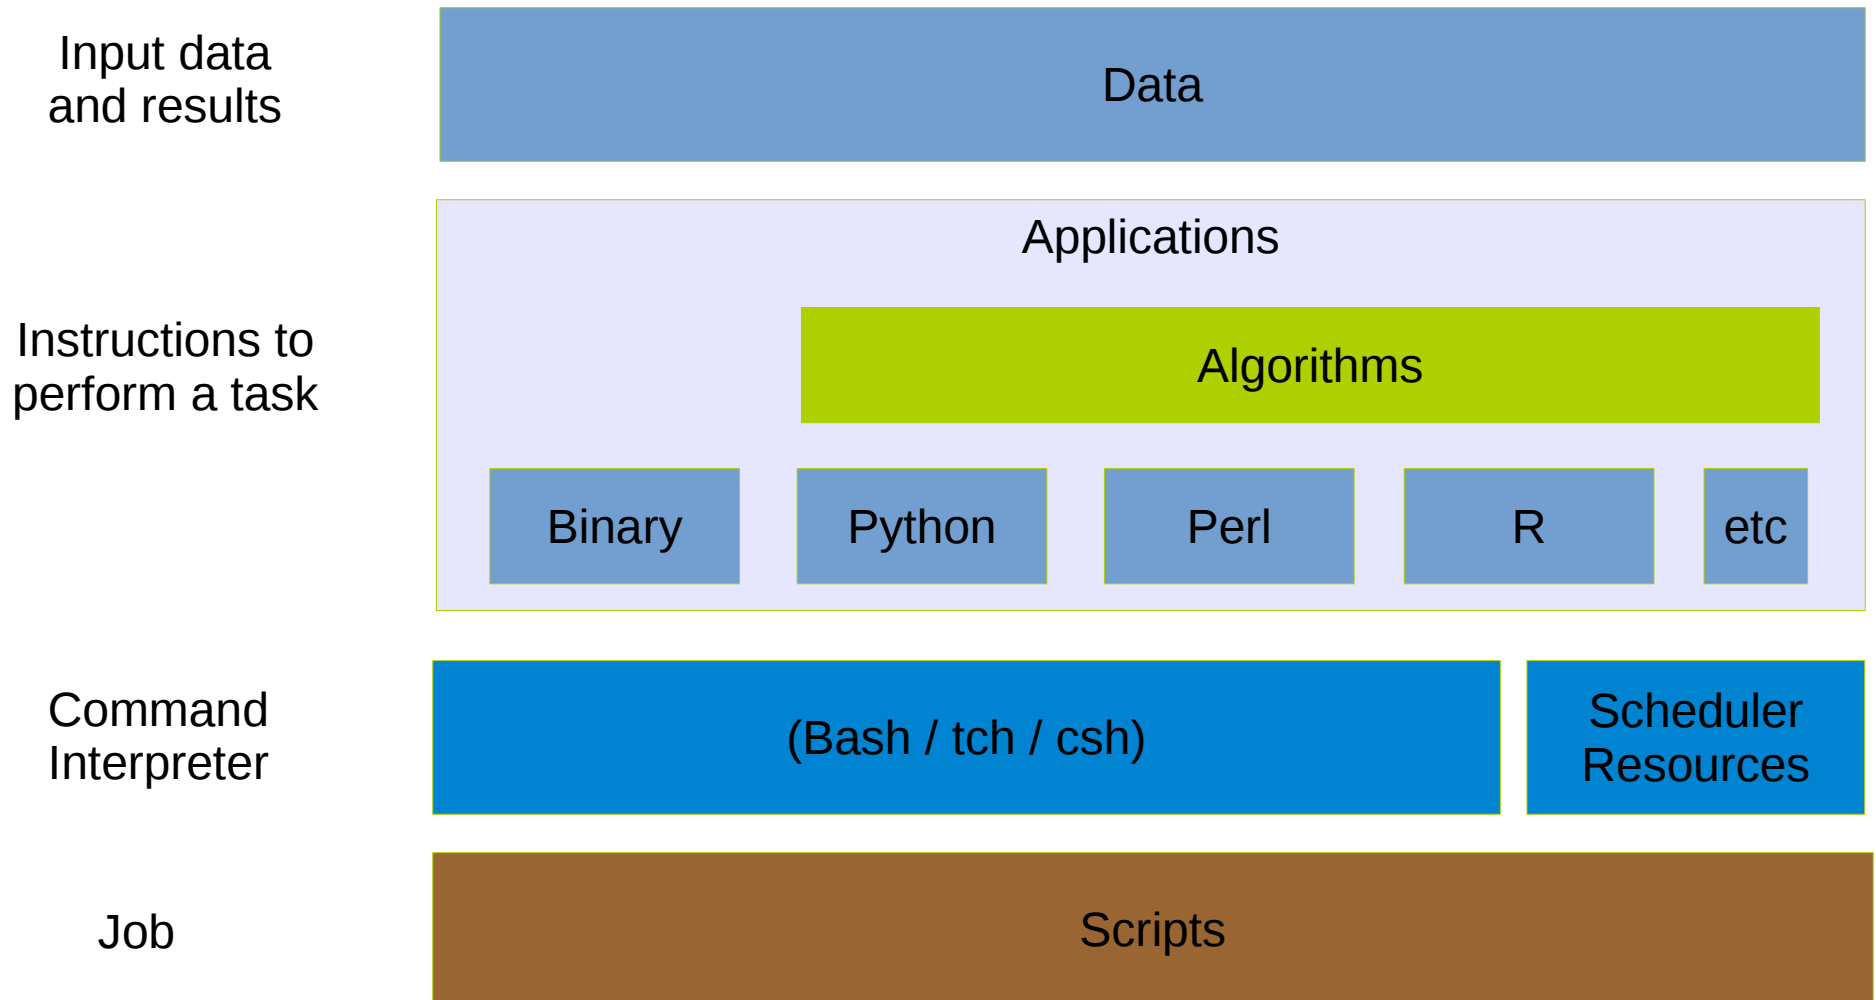
# Jobs: Parallel v/s Distributed



# Working with a HPC System



# Job Scripting



# Job Scheduler Directives

```
#!/bin/bash
# Resource specification
#$ -l h_rt=1:00:00
#$ -cwd
#$ -j y
#$ -V
#$ -notify
# User Notification
#$ -m abes
#$ -M myemail@domain.com
# Job name
#$ -N jobname
# Command interpreter
#$ -S /bin/bash
# Parallel environment: openmpi,openmp,etc
#$ -pe openmpi 128
# Job Array
#$ -te 1:1000
# Queue to use
#$ -q all.q
```

- **Grid Engine**
- PBS
- Condor

# Job Scheduler Directives

```
#!/bin/bash
# number of nodes and processes per node
#PBS -l select=4:mpiprocs=8
# resources
#PBS -l mem=213mb
#PBS -l walltime=2:00:00
#PBS -l cput=1:00:00
# name of job
#PBS -N jobname
# User notificacion
#PBS -m bea
#PBS -M myemail@domain.com
# Use submission environment
#PBS -V
# Queue to use
#PBS -q default
```

- Grid Engine
- **PBS**
- Condor

# Job Scheduler Directives

```
## executable file
Executable = my_program.exe
## Specify Condor execution environment.
Universe = vanilla
## execution machines
Requirements = ((OpSys == "LINUX"))
## Define input files and arguments
Input = stdin.txt.$(Process)
Arguments = in.$(Process) out.$(Process)
## Define output/error/log files
Output = logs/stdout.$(Process).txt
Error = logs/stderr.$(Process).txt
Log = logs/log.$(Process).txt
## files need to be transferred and when.
Transfer_input_files = files/in1.$(Process)
Transfer_output_files = out.$(Process)
Transfer_executable = true
Should_transfer_files = YES
When_to_transfer_output = ON_EXIT
## Add 100 copies of the job to the queue
Queue 100
```

- Grid Engine
- PBS
- **Condor**

# Environment Modules

- Configure the environment to run a particular application (or a set of applications)
  - Environmental variables:
    - PATH
    - LD\_LIBRARY\_PATH
    - LD\_RUN\_PATH}
  - Library versions and locations
    - BOOST\_HOME, ATLAS\_HOME, etc
  - Compilation & execution flags
    - CFLAGS, LDFLAGS, CXXFLAGS, etc.

# Environment Modules

- Example: module available

```
[astrouser@levque ~]$ module available
```

```
----- /usr/share/Modules/modulefiles -----  
dot          module-cvs  module-info  modules      null         use.own
```

```
----- /etc/modulefiles -----  
R/3.1.0      intel/12.1.5  openmpi_gcc44/1.4.3  
adf/2012.01  intelmpi/4.0.2  openmpi_intel/1.4.3  
astro/1.1    intelmpi/4.0.3  openmpi_intel/1.6.5  
cplex/12.1.0  lsst/9.2       openmpi_pgi/1.4.3  
cplex/12.2.0  namd/2.7       opl/6.3  
cplex/9.1.0  namd/2.8       pgi/10.9  
espresso/5.1  namd/2.9       python/2.6.6  
gaussian/09B1  namd_intel/2.8  siesta/3.2  
gnuplot/4.6.5  octave/3.8.1   stata/11.0  
gromacs/4.6.5  openmpi/1.4.2  
gromacs_single/4.5.3  openmpi/1.4.3
```

```
[astrouser@levque ~]$
```

# Environment Modules

- `module show {module name/version}`

```
[astrouser@levque ~]$ module show lsst
```

```
-----  
/etc/modulefiles/lsst/9.2:
```

```
module-whatis      Sets up the LSST Astro toolchain (Self-standing edition)  
                   in you enviornment  
prepend-path       PATH /home/apps/lsst/bin  
prepend-path       PATH /home/apps/lsst/sbin  
prepend-path       LD_LIBRARY_PATH /home/apps/lsst/lib  
prepend-path       LD_LIBRARY_PATH /home/apps/lsst/lib64  
prepend-path       MANPATH /home/apps/lsst/share/man
```

```
-----  
[astrouser@levque ~]$
```

# Environment Modules

- module load {module name/version}
- module list

```
[student01@syntagma ~]$ module load lsst
```

```
[student01@syntagma ~]$ module list
```

```
Currently Loaded Modulefiles:
```

```
 1) lsst/9.2
```

```
[astrouser@development ~]$ echo $LD_LIBRARY_PATH
```


```
/home/apps/lsst/lib64:/home/apps/lsst/lib:/opt/intel/Compiler/11.1/072/lib/intel64:  
/opt/intel/Compiler/11.1/072/ipp/em64t/sharedlib:/opt/intel/Compiler/11.1/072/mkl/lib  
/em64t:/opt/intel/Compiler/11.1/072/lib/intel64:/opt/intel/Compiler/11.1/072/ipp/em64t  
/sharedlib:/opt/intel/Compiler/11.1/072/mkl/lib/em64t
```

```
[astrouser@development ~]$ echo $PATH
```

```
/home/apps/lsst/sbin:/home/apps/lsst/bin:/opt/gridengine/bin/lx24-amd64:/usr/lib64  
/qt-3.3/bin:/usr/kerberos/bin:/opt/intel/Compiler/11.1/072/bin/intel64:/opt/intel  
/Compiler/11.1/072/bin/intel64:/usr/local/bin:/bin:/usr/bin:/home/uchile/cmm  
/astrolab/astrouser/bin
```



# Creating (Sun Grid Engine) Jobs



```
#!/bin/bash
#$ -cwd
#$ -j y
#$ -V
#$ -notify
#$ -m abes
#$ -M myemail@domain.com
#$ -N my-first-job
#$ -S /bin/bash
#$ -q all.q

module load lsst

echo "Running at `hostname -s`"
echo "Starting at `date +%c`"

INPUT_FITS=$1
WEIGHT_FITS=$2

sex $INPUT_FITS -CATALOG_NAME catalogue.cat \
    -WEIGHT_IMAGE $WEIGHT_FITS

echo "done"
```

# Interacting with the SGE

- `qsub job_script.sge`
  - Submit *job\_script* to the queue
- `qstat`
  - Show only the status of **your jobs** in the queue
- `qstat -f`
  - Show the status of **jobs** by queue
- `qstat -u "user"`
  - Show the jobs owned by `user`
- `qstat -j Job-ID`
  - Show the status for the job `Job-ID`
- `qhost`
  - Show the status of each node
- `qhost -j`
  - Show the status of each node and the jobs running on each one of them
- `qdel Job-ID`
  - Cancel (running) and delete a job from the queue

# Submitting & Monitoring Jobs

```
[astrouser@levque ~]$ qsub my_first_job.sge
Your job XXXX ("my-first-job") has been submitted
```

```
[astrouser@levque ~]$ qstat
job-ID  prior   name           user    state submit/start at   queue                          slots ja-task-ID
-----
XXXX 0.00000 my-first-s    astrouser  qw    08/09/2013 23:00:06                          1
```

```
[astrouser@levque ~]$ qstat
job-ID  prior   name           user    state submit/start at   queue                          slots ja-task-ID
-----
XXXX 0.00000 my-first-s    astrouser  r     08/09/2013 23:00:30 all.q@levque001  1
```

- watch is your friend
  - watch -n 1 "qstat" : show qstat at 1 second interval
- Ganglia is your best friend

# Monitoring Jobs

- **Ganglia** is an open source monitoring system developed in the NPACI (UCLA) and used initially by the UCSD to monitor their Rocks clusters.

<http://syntagma.cmm.uchile.cl/ganglia>

<http://levque.dim.uchile.cl/ganglia>

- Queue is monitored at “host overview” in the frontend.
- Compute nodes “host overview” gives you the state of your processes (require an extra plug-in)
- Useful metrics such as memory and network consumption are shown in an aggregated way as well as in a host basis way.

# Best Practices

# Rule 1 : Characterize your job

Because you are not alone in the world,  
when submitting your job, please mind:

- The **memory** consumption.
  - Per core
  - Per process
- How **intensive** is on **Input/Output** operations.
  - Is my process able to run in a NFS storage layout?
  - Is my process able to run in a parallel filesystem
  - Do my process need a locking scheme to work?
- Scalability in terms of **cores**.
  - No always the highest, the better

# Rule 2 : Monitor your Job

Because your jobs are running together with other people's job, please mind:

- Check continuously the status of the compute nodes where your jobs are running
- Check the load of the filesystem if your job is I/O intensive
  - Ganglia is your friend
- Be sure about how much storage your job is consuming.
  - You don't want to be remembered as the guy who crash everyone else jobs due to “out of space” in the filesystem.

# Rule 3 : Do not leave data unattended

Because the storage is expensive and other users have the same right to use it, please mind:

- Remove temporal data and job output files.
  - .o and .e files
  - By applying the Rule 1, you should know if your job produces temporal files that can be deleted at the end of the job.
- When finishing with your results, keep them in other storage than your HPC account.
- Do not leave data in the scratch filesystem – It is volatile, which means it can be wiped out without further notice.
- Get used to backup your work (scripts, jobs, results).



# Rule 4: Share your knowledge

Because not all the users have your skills to work with a HPC System, please mind:

- Write documentation by commenting your jobs scripts between lines.
- Write `readme` files specifying how many “*shaolin chambers*” did you have to go through to get your job done.
- Publish your scripts (if you are allowed to)

# Hands-On

“My first sextraction”

# Accessing the HPC System

SSH Connection



Ssh keys generation  
Or Password change  
(only first time)



Login OK

```
# ssh astrouser@development.cmm.uchile.cl
```

```
astrouser@development.cmm.uchile.cl's password:  
Last login: Sat Aug 16 19:08:44 2014 from pc-163-125-  
44-190.cm.vtr.net
```

```
Bienvenido a Levque
```

```
...
```

```
[astrouser@development ~]$ ssh-keygen -t rsa
```

```
[astrouser@development ~]$ passwd
```

```
...
```

```
...
```


```
[astrouser@development ~]$
```

# Accessing the HPC System

Checking



Home directory listing




Ready to go!

```
[astrouser@development ~]$ ls -la
total 692
drwx----- 26 astrouser astro 4096 ago 16 19:08 .
drwxr-xr-x 17 root root 4096 ago 14 20:07 ..
lrwxrwxrwx 1 astrouser astro 16 jun 27 12:30 astro ->
/home/apps/astro
drwxr-xr-x 4 astrouser astro 4096 jul 2 18:15 .astropy
-rw----- 1 astrouser astro 12936 ago 15 15:14 .bash_history
-rw-r--r-- 1 astrouser astro 33 jun 25 16:06 .bash_logout
-rw----- 1 astrouser astro 42 ago 14 17:14 .lessht
lrwxrwxrwx 1 astrouser astro 16 ago 13 10:54 lsst ->
/home/apps/lsst/
drwxr-xr-x 3 astrouser astro 4096 ago 14 19:04 .matplotlib
drwxr-xr-x 2 astrouser astro 4096 ago 14 19:54 .mc
drwxr-xr-x 5 astrouser astro 4096 jul 12 10:08 .mozilla
-rw-r--r-- 1 astrouser astro 4322 jul 7 16:26 .recently-used.xbel
drwx----- 2 astrouser astro 4096 ago 13 19:09 .ssh
drwxr-xr-x 3 astrouser astro 4096 jun 27 12:46 .subversion
-rw-r--r-- 1 astrouser astro 658 jun 25 16:06 .zshrc
[astrouser@development ~]$
```


# Getting Cluster State

## Checking Hosts



```
[astrouser@development ~]$ qhost
HOSTNAME          ARCH          NCPU  LOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS
-----
global            -             -     -     -       -       -       -
levque001         lx24-amd64    8     0.02  23.5G   1.2G    2.0G    64.4M
levque002         lx24-amd64    8     5.00  23.5G   1.9G    2.0G    50.0M
...
levque065         lx24-amd64    8     7.00  23.5G   7.8G    2.0G    88.0K
levque066         lx24-amd64    8     0.00  23.5G   593.3M  2.0G    6.6M
matlab1           lx24-amd64   16     0.10  23.5G   493.4M  2.0G     0.0
```

## Checking Jobs



```
[astrouser@development ~]$ qstat
job-ID prior  name          user          state submit/start at      queue          slots
ja-task-ID
-----
1239018 0.99951 CR_Blind14    astrouser     r      08/12/2014 11:50:27 all.q@levque062      4 30
1239021 0.00000 doastro_Bl    astrouser     hqw    08/06/2014 23:06:30                4 1
1239022 0.00000 doastro_Bl    astrouser     hqw    08/06/2014 23:06:30                4 1
1239023 0.00000 doastro_Bl    astrouser     hqw    08/06/2014 23:06:30                4 1
```


## Ready to go!

```
[astrouser@development ~]$
```

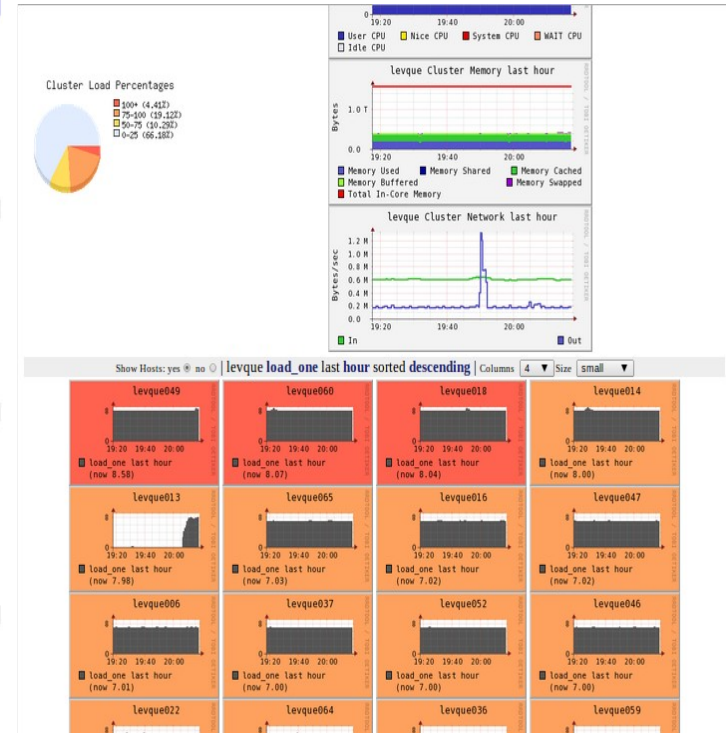
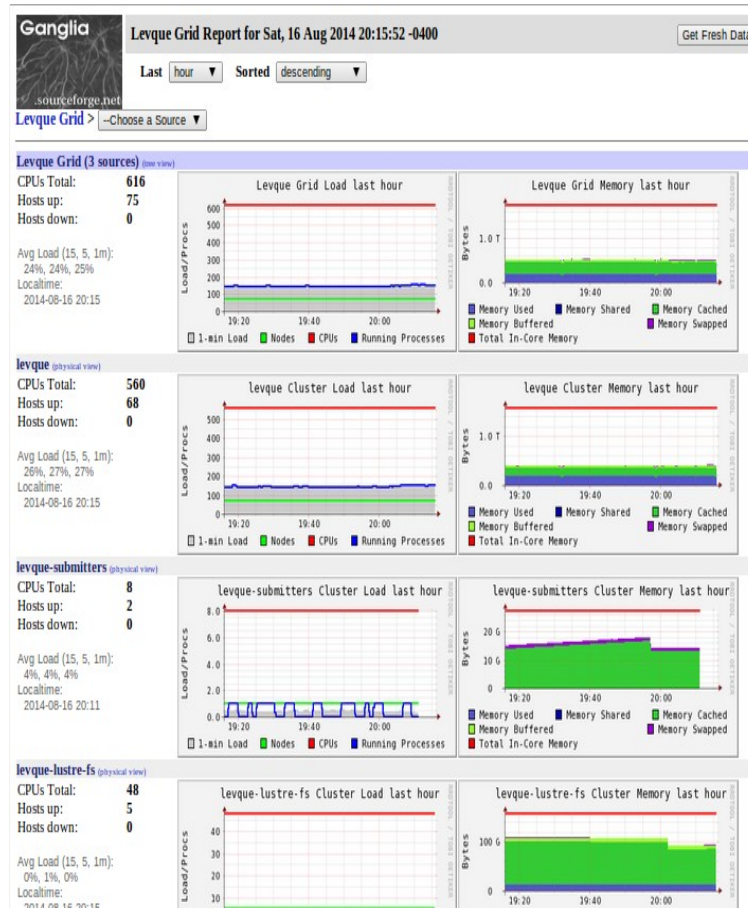
# Getting Cluster State

<http://levque.cmm.uchile.cl/ganglia>

Via web



Ready to go!



# Preparing My Data

Cloning base files

```
[astrouser@levque demos]$ git clone http://git.cmm.uchile.cl/demo-sextractor.git  
git clone http://git.cmm.uchile.cl/demo-sextractor.git
```

Create working directory

```
[astrouser@levque demos]$ mkdir -p workdir
```

Copy input files

```
[astrouser@levque demos]$ cp demo-sextractor/Blind_03_N1_01* workdir
```

Copy sextractor  
default files & job script

```
[astrouser@levque demos]$ cp demo-sextractor/default.* workdir/
```

```
[astrouser@levque demos]$ cp demo-sextractor/*.sge workdir/
```

Getting into the cwd


```
[astrouser@levque demos]$ cd workdir
```

Ready to go!

```
[astrouser@levque demos/workdir]$
```

# Preparing My Job

Edit the job script



```
[astrouser@levque demos/workdir]$ cat my-first-sextraction.sge
#!/bin/bash
# My First Sextraction (by using sextractor)
# usage: my-first-sextraction.sge {input_fits} {weight_fits}
#
# By Francisco Forster and Juan-Carlos Maureira
# CMM. University of Chile
#####
#$ -cwd
#$ -j y
#$ -notify
#$ -m abes
#$ -M your-email@mail.com
#$ -N my-first-sextraction
#$ -S /bin/bash

# load the proper enviroment module
module load astro

# check for arguments
if [ $# != 2 ]; then
    echo "usage: $0 {input_fits} {weight_fits}"
    Exit 1
Fi

# get the input arguments (see usage)
INPUT_FITS=$1
WEIGHT_FITS=$2

if [ ! -f "$INPUT_FITS" ] || [ ! -f "$WEIGHT_FITS" ]; then
    echo "input files not found."
    Exit 1
Fi

# perform the source extraction and generate the file catalogue.cat
sex $INPUT_FITS -CATALOG_NAME catalogue.cat -WEIGHT_IMAGE $WEIGHT_FITS

echo "done"
[astrouser@levque demos/workdir]$
```

Ready to go!



# Characterizing My Job

Open two terminals

1: execute the  
getProcessMemoryStatus

2: execute the script by hand

```
[astrouser@levque demos/workdir]$ ./getProcessMemoryStatus sex
Sample VmPeak: VmSize: VmLck: VmHWM: VmRSS: VmData: VmStk: VmExe: VmLib:
VmPTE: VmSwap:
1 81220 81216 0 12116 12116 22080 92 1368 37268 124 0
2 81220 81216 0 12120 12112 22080 92 1368 37268 124 0
3 81220 81212 0 12124 12120 22076 92 1368 37268 124 0
4 81220 81212 0 12124 12116 22076 92 1368 37268 124 0
5 96908 96908 0 23128 23128 37772 92 1368 37268 132 0
6 105392 105392 0 37296 37296 46256 92 1368 37268 136 0
7 106408 105612 0 38016 37432 46476 92 1368 37268 136 0
8 106408 105612 0 38016 37456 46476 92 1368 37268 136 0
9 106408 105612 0 38016 37460 46476 92 1368 37268 136 0
10 106408 105612 0 38016 37548 46476 92 1368 37268 136 0
11 106408 105944 0 38016 37788 46808 92 1368 37268 136 0
12 114240 105468 0 45796 37512 46332 92 1368 37268 136 0
13 114240 105624 0 45796 37556 46488 92 1368 37268 136 0
14 114240 105624 0 45796 37564 46488 92 1368 37268 136 0
15 114240 78328 0 45796 11304 19192 92 1368 37268 116 0
^C
[astrouser@levque demos/workdir]$
```

```
[astrouser@levque demos/workdir]$ ./my-first-
sextraction.sge \
    Blind_03_N1_01.fits.fz_proj.fits \
    Blind_03_N1_01_wtmap.fits.fz_proj.fits
Running at development.cmm.uchile.cl
Starting time : Sat 10 Aug 2013 12:41:00 AM CLT
.
----- SExtractor 2.18.9 started on 2013-08-10 at
00:41:00 with 1 thread

Setting catalog parameters.
Reading detection filter.
Initializing catalog.
Looking for Blind_03_N1_01.fits.fz_proj.fits.
----- Measuring from: Blind_03_N1_01.fits.fz_p
...

[astrouser@levque demos/wordir]$
```

# Submitting My Job

```
[astrouser@levque demos/workdir]$ qsub my-first-sextraction.sge \  
Blind_03_N1_01.fits.fz_proj.fits \  
Blind_03_N1_01_wtmap.fits.fz_proj.fits
```

Your job 9762 ("my-first-sextraction") has been submitted

```
[astrouser@levque demos/workdir]$ qstat
```

job-ID	prior	name	user	state	submit/start at	queue	slots	ja-task-ID
9762	0.00000	my-first-s	astrouser	qw	08/09/2013 23:00:06		1	

```
[astrouser@levque demos/workdir]$ qstat
```

job-ID	prior	name	user	state	submit/start at	queue	slots	ja-task-ID
9762	0.00000	my-first-s	astrouser	r	08/09/2013 23:00:16	all.q@levque001	1	

```
[astrouser@levque demos/workdir]$
```

Submit the job



Check submission status

You got an Email!

# Monitoring My Job

```
[astrouser@levque demos/workdir]$ watch -n 1 'qstat'
```

```
Every 1.0s: qstat
```

job-ID	prior	name	user	state	submit/start at	queue	slots	ja-task-ID
9762	0.00000	my-first-s	astrouser	qw	08/09/2013 23:00:06		1	

My job only

```
Fri Aug 9 23:12:52 2013
```

```
[astrouser@levque demos/workdir]$ watch -n 1 "qstat -u '*'"
```

```
Every 1.0s: qstat -u "*"
```

job-ID	prior	name	user	state	submit/start at	queue	slots	ja-task-ID
9051	0.50500	QLOGIN	dtravisany	r	07/09/2013 16:14:16	all.q@levque003	1	
9683	0.50500	QLOGIN	adigenova	r	07/28/2013 18:21:19	all.q@levque034	1	
9703	0.60500	V6O18H12_s	wcanon	r	08/02/2013 13:44:15	all.q@levque023	8	
9704	0.60500	V6O18H12_s	wcanon	r	08/02/2013 13:51:45	all.q@levque056	8	
9741	0.50500	QLOGIN	mcortes	r	08/08/2013 15:18:02	all.q@levque043	1	
9759	0.60500	BIphenFeCs	wcanon	r	08/09/2013 19:52:31	all.q@levque065	8	
9761	0.60500	BIphenFeC2	wcanon	r	08/09/2013 20:28:16	all.q@levque042	8	
9762	0.00000	my-first-s	astrouser	r	08/09/2013 23:00:16	astro.q@levque001	1	
9760	0.60500	BIphenFeC2	wcanon	qw	08/09/2013 20:10:11		8	

All the queue

```
Fri Aug 9 23:13:12 2013
```

```
[astrouser@levque demos/workdir]$
```

# Getting My Job Status

```
[astrouser@levque demos/workdir]$ qstat -j 9762
```

```
=====
job_number:                9762
submission_time:           Fri Aug  9 23:05:29 2013
owner:                     astrouser
Uid:                       18010
group:                     demo
Gid:                       11056
sge_o_home:                /home/uchile/cmm/astrolab/astrouser
sge_o_log_name:            astrouser
sge_o_shell:               /bin/bash
sge_o_workdir:             /home/uchile/cmm/astrolab/astrouser/demos/workdir
sge_o_host:                login-1-0
account:                   sge
cwd:                       /home/uchile/cmm/astrolab/astrouser/demos/workdir
merge:                     y
mail_options:              abes
mail_list:                 jcm@dim.uchile.cl
notify:                    TRUE
job_name:                  my-first-sextraction
Jobshare:                  0
shell_list:                NONE:/bin/bash
error reason    1:         08/09/2014 01:07:49 [15086:29399]: error: can't open output file "/home/apps/astro/logs/XXXXXXX.o
env_list:
job_args:                 Blind_03_N1_01.fits.fz_proj.fits,Blind_03_N1_01_wtmap.fits.fz_proj.fits
scheduling info:         queue instance "all.q@levque001" dropped because it is temporarily not available
                        queue instance "all.q@levque002" dropped because it is temporarily not available
                        queue instance "all.q@levque003" dropped because it is disabled
                        queue instance "all.q@levque004" dropped because it is disabled
                        queue instance "all.q@levque065" dropped because it is disabled
                        queue instance "all.q@levque066" dropped because it is full
                        queue instance "all.q@matlab1" dropped because it is full
                        job dropped because of job dependencies
```

Job information

Why the job is not  
In that compute node


```
[astrouser@levque demos/workdir]$
```

# Getting My Results

Looking for result files



Checking results contents



YES!!!

```
[astrouser@levque demos/workdir]$ ls -la
total 119732]
drwxr-xr-x 3 astrouser astro      4096 ago 16 19:13 .
drwxr-xr-x 3 astrouser astro      4096 ago 16 19:10 ..
-rw-r--r-- 1 astrouser astro 61228800 ago 16 19:10 Blind_03_N1_01.fits.fz_proj.fits
-rw-r--r-- 1 astrouser astro      40958 ago 16 19:10 Blind_03_N1_01.fits.fz_proj.fits-
catalogue.dat
-rw-r--r-- 1 astrouser astro 61228800 ago 16 19:10 Blind_03_N1_01_wtmap.fits.fz_proj.fits
-rw-r--r-- 1 astrouser astro      42702 ago 16 19:11 catalogue.cat
-rw-r--r-- 1 astrouser astro        804 ago 16 19:10 default.conv
-rw-r--r-- 1 astrouser astro         82 ago 16 19:10 default.param
-rw-r--r-- 1 astrouser astro     11520 ago 16 19:10 default.psf
-rw-r--r-- 1 astrouser astro      3620 ago 16 19:10 default.sex
-rwxr-xr-x 1 astrouser astro        865 ago 16 19:10 getProcessMemoryStatus
drwxr-xr-x 8 astrouser astro      4096 ago 16 19:13 .git
-rwxr-xr-x 1 astrouser astro     1042 ago 16 19:10 my-first-sextraction.sge
-rw-r--r-- 1 astrouser astro        711 ago 16 19:10 README.seaf
-rw-r--r-- 1 astrouser astro      3620 ago 16 19:10 my-first-sextraction.o9762
```

```
[astrouser@levque demos/workdir]$ head -n 5 my-first-sextraction.o9762
Running at levque001
Starting time : Fri 09 Aug 2013 11:00:16 PM CLT
.
----- SExtractor 2.18.9 started on 2013-08-09 at 23:24:46 with 1 thread
```

```
[astrouser@levque demos/workdir]$ tail -n 5 my-first-sextraction.o9762
 433 1477.1 2924.3   3.4  3.2      115090 _____
 434 1914.5 3765.5   2.3  2.1       646.094 _____
 435  806.2 2865.6   2.9  2.6       2201.14 _____
 436 1699.9 2835.1   2.3  2.0       522.624 _____
Objects: detected 443      /  sextracted 436
```


```
Closing files.
.
Ending time : Fri 09 Aug 2013 11:24:47 PM CLT
done
```

```
[astrouser@levque demos/workdir]$
```


# Hands-on: My first job in HPC@CMM

## Saving and Cleaning up

Return to your \$HOME



Save results files



Clean the working  
directory



logout

```
[astrouser@levque demos/workdir]$ cd
```

```
[astrouser@levque demos]$ cp workdir/cataloge.cat demo-sextractor
```

```
[astrouser@levque demos]$ rm -rf workdir
```

```
[astrouser@levque demos]$ logout
```

# The Take Aways

- Definitions needed to understand a HPC system.
- Overview about architecture and components of a HPC system.
- Software, Applications, tools-chains
- Best practices when working with a HPC system.
- Follow an example of how to work with generic HPC system.