

Working on the Irene-amd machine

Table of Content

Working on the Irene-amd machine	1
1. Connexion	2
2. Introduction	2
3. Job manager commands	2
4. Suggested environment	2
4.1. General environment	2
4.2. Subversion version	3
5. How the storage project is set by libGCM	3
6. File systems used on Irene	4
7. Irene job headers	4
8. Use of specific options to increase computing performances	5
8.1. Increase the number of cores per MPI/OpenMP process	5
8.2. Use of dedicated nodes for XIOS servers	5

1. Connexion

```
ssh login@irene-amd-fr.ccc.cea.fr
```

2. Introduction

On-line users manual: <https://www-tgcc.ccc.cea.fr/docs/irene.info.pdf> (you will need a TGCC login and password)

- Partition name : Rome
- CPUs: 2x64 AMD Rome@2.6Ghz (AVX2)
- Core/Node: 128
- Nodes: 2292
- Total core: 293376
- RAM/Node: 256GB
- RAM/core : 2GB
- When submitting a job through `ccc_msub` or `ccc_mprun`, you must specify `-m work`, `-m scratch`, `-m store`, or combine them like in `-m work,scratch`. This constraint has the advantage that your jobs won't be suspended if a file system you don't need becomes unavailable. This is done in all jobs in libIGCM.
- The default time limit for a job submission is 2 hours (7 200 s).
- Irene post-processing nodes : `xlarge` are free and useful for post-processing operations.
 - Fat nodes for computation requiring a lot of shared memory (coming from irene.info)
 - Partition name: `xlarge`
 - CPUs: 4x28-cores Intel Skylake@2.1GHz
 - GPUs: 1x Nvidia Pascal P100
 - Cores/Node: 112
 - Nodes: 5
 - Total cores: 560
 - RAM/Node: 3TB
 - RAM/Core: 5.3GB
 - IO: 2 HDD de 1 TB + 1 SSD 1600 GB/NVMe

3. Job manager commands

- `ccc_msub job ->` submit a job
- `ccc_mdcl ID ->` kill the job with the specified ID number
- `ccc_mstat -u login ->` display all jobs submitted by login, add `-f` to see full job name
- `ccc_mpp ->` display all jobs submitted on the machine. `ccc_mpp -n` to avoid colors.
- `ccc_mpp -u $(whoami) ->`display your jobs.

4. Suggested environment

4.1. General environment

Before working on Irene you need to prepare your environment. This is important to do before compilation to ensure the use of same modules as done by libIGCM running environment. We propose you 2 files which you can copy from the home igcmg. The first one called `bashrc` will source the second called `bashrc_irene`. Copy both files to your home, rename them by adding a dot as prefix. You can add personal settings in your `.bashrc_irene-amd`. Do as follow:

```
cp ~igcmg/MachineEnvironment/irene-amd/bashrc ~/.bashrc # copy it even if you already have a .bashrc file
cp ~igcmg/MachineEnvironment/irene-amd/bashrc_irene-amd ~/.bashrc_irene-amd
```

The `.bashrc` will source your own `.bashrc_irene-amd` which must be in your home.

After re-connexion or source of `.bashrc_irene-amd`, check your loaded modules for intel, netcdf, mpi, hdf5 needed for the compilation:

```

module list
Currently Loaded Modulefiles:
 1) ccc/1.0(default)          6) feature/mkl/lp64          11) c/intel/19.0.5.281(default)  16) feature/openmpi/
 2) datadir/drf              7) feature/mkl/sequential    12) c++/intel/19.0.5.281(default) 17) flavor/openmpi/s
 3) datadir/own(default)     8) feature/mkl/single_node   13) fortran/intel/19.0.5.281(default) 18) feature/openmpi/
 4) dfldatadir/own(default)  9) mkl/19.0.5.281           14) intel/19.0.5.281(default)     19) .tuning/openmpi/
 5) flavor/buildcompiler/intel/19 10) licsrv/intel             15) flavor/buildmpi/openmpi/4.0    20) hwloc/2.0.4

```

The modules are specified in the file `~igcmg/MachineEnvironment/irene-amd/env_irene-amd` which is sourced in `bashrc_irene-amd`. The same file `env_irene-amd` is sourced in `libIGCM`.

--> Be careful this environment can be update during next weeks according to TGCC recommendations

Create `~/.forward` file in your main home containing only one line with your email address to receive emails from `libIGCM`.

4.2. Subversion version

Since only recent subversion version (i.e > 1.6) are installed on Irene supercomputer, some usual functionalities are not available anymore (ex : `svn` command on copy of subdirectories...). In order to keep these functionalities, subversion 1.6.9 has been installed. To use this version :

```

module unload subversion
module use ~igcmg/Modules/tools
module load subversion/1.6.9
irenel90 : svn --version

svn, version 1.6.9 (r901367)

compiled Oct 31 2018, 11:12:49

```

The use of this version has been added in default environment `~igcmg/MachineEnvironment/irene-amd/bashrc_irene-amd`. Beware of coherence of subversion version you use from the extraction of your model/configuration to the use of `svn` commands in directories of your model/configuration.

5. How the storage project is set by libIGCM

When you use `libIGCM` it is recommended to dedicate one `modipsl/libIGCM` to one project allocation. By default, the output folders `IGCM_OUT` will be created in the directories `$CCCSCRATCHDIR`, `$CCCWORKDIR` and `$CCCSTOREDIR` corresponding to the project used in the main job. It is important that the same project is used in the post-processing jobs in `libIGCM`.

For `gencmip6` project, you have to set a subproject for computing and `gencmip6` is forced for all directories.

If you need to use another project for the computing than the storage, it is possible to set the variable **DataProject** in `config.card` `UserChoices` section, for example `DataProject=gen6328`, read more [here](#). This project will be used for all output directories for the computing job and post-processing jobs even if they have another project for computing in the headers. The variable **DataProject** can also be used if you work with different project allocations in the same `modipsl`. Only exception (harmless) is the first `RUN_DIR` folder which is always created in the `$CCCSCRATCHDIR` corresponding to the `dfldatadir` loaded in when submitting main job. When the job resubmits itself, the `RUN_DIR` will be in the same project space as the rest of the output. #

Example of `ins_job`

```

> ./ins_job
...
Wait for the next question ...
Hit Enter or give project ID (default is gen0000), possible projects are gen1111 gen2222 ... or other xxxcmip6 : aaacmip6
gen0000 (RETURN)
ProjectID is gen0000 at Irene
Hit Enter or give TYPE OF NODE required for post-processing (default is "xlarge"), possible types of nodes are "skylake" o
(RETURN)
ProjectNode for post-processing is xlarge at Irene
Hit Enter or give NUMBER OF CORES required for post-processing (default is "8")
possible numbers of cores are "1" to "112" for xlarge :
(RETURN)
ProjectCore for post-processing is 8

```

```

Wait for the next question ...
Hit Enter or give project ID (default is gen0000), possible projects are gen1111 gen2222 ... or other xxxcmip6 : aaacmip6
(RETURN)
PostID is gen0000 at Irene on xlarge for post-processing

```

6. File systems used on Irene

A figure to illustrate Irene filesystems is available [here](#)

7. Irene job headers

Here is an example of a job header as generated by libGCM on the Irene-amd machine:

```

#####
## IRENE  TGCC/CEA ##
#####
#MSUB -r MY-SIMULATION
#MSUB -o Script_Output_MY-SIMULATION.000001
#MSUB -e Script_Output_MY-SIMULATION.000001
#MSUB -eo
#MSUB -n 976
#MSUB -x
#MSUB -T 86400
#MSUB -A dekc mip6
#MSUB -q rome
#MSUB -m store,work,scratch

```

The detail is as follows:

Control	Keyword	Argument	Example	Comments
<i>Job name</i>	-r	string	#MSUB -r Job_MY-SIMULATION	The string should not contain the underscore (_) character.
<i>Standard output file name</i>	-o	string	#MSUB -o Script_Output_MY-SIMULATION.000001	
<i>Error output file name</i>	-e	string	#MSUB -e Script_Output_MY-SIMULATION.000001	If both -o and -e names are the same, the two outputs will be merged. The -eo option of the example is as a matter of fact redundant with that.
<i>Number of MPI tasks allocated</i>	-n	integer	#MSUB -n 976	This is for the SPMD case; in the MPMD case, the integer is the number of cores.
<i>Node exclusivity</i>	-x	none	#MSUB -x	
<i>Wall-time (maximum time allowed for execution)</i>	-T	integer	#MSUB -T 86400	Time is expressed in number of seconds.
<i>Project allocation</i>	-A	string	#MSUB -A dekc mip6	
<i>Partition used</i>	-q	string	#MSUB -q rome	Choice is amongst rome

<i>Visible spaces</i>	-m	(comma-separated) string(s)	#MSUB -m store,work,scratch	Specific to TGCC. The job can only "see" spaces specified on this line. Possible spaces are amongst store and/or work and/or scratch.
<i>Job priority</i>	-U	string	#MSUB -U high	Possible arguments are high, medium or low. (This command does not appear by default on jobs generated by libIGCM.)

8. Use of specific options to increase computing performances

Following options have to be used with caution. Please contact the person in charge of the configuration you use before launching production runs with these options.

8.1. Increase the number of cores per MPI/OpenMP process

Because of the architecture of Irene-amd processor, you could be limited by the memory available per core as well as the memory bandwidth. In order to use more cores per MPI process (or OpenMP thread), you can specify the factor you want to increase this number of cores in config.card by using **DEP** keyword as follows :

```
ATM= (gcm.e, lmdz.x, 71MPI, 8OMP, 2DEP)
SRF= ( " " , " " )
SBG= ( " " , " " )
OCE= (opa, opa.xx, 360MPI, 2DEP)
ICE= ( " " , " " )
MBG= ( " " , " " )
CPL= ( " " , " " )
IOS= (xios_server.exe, xios.x, 12MPI, 2DEP)
```

In this example, the number of cores will be increased by a factor 2 for each component.

8.2. Use of dedicated nodes for XIOS servers

Some preliminary tests (bad performances) drive us to investigate the use of dedicated nodes for XIOS servers. To activate this functionality, you have to modify your config.card by using "DEDICATED" keyword as follows :

```
ATM= (gcm.e, lmdz.x, 71MPI, 8OMP, 2DEP)
SRF= ( " " , " " )
SBG= ( " " , " " )
OCE= (opa, opa.xx, 360MPI, 2DEP)
ICE= ( " " , " " )
MBG= ( " " , " " )
CPL= ( " " , " " )
IOS= (xios_server.exe, xios.x, 12MPI, 2DEP, DEDICATED)
```

Note that this functionality is only available for IOS component.