

## Working on the Irene machine

---

### Table of Content

<b>Working on the Irene machine</b>	<b>1</b>
<b>1. Introduction</b>	<b>2</b>
<b>2. Job manager commands</b>	<b>2</b>
<b>3. Suggested environment</b>	<b>2</b>
3.1. General environment	2
3.2. Subversion version	3
<b>4. File system</b>	<b>3</b>
<b>5. How the storage project is set by libIGCM</b>	<b>4</b>
<b>6. A figure to understand file systems used on Irene</b>	<b>4</b>

## 1. Introduction

- On-line users manual: <https://www-tgcc.ccc.cea.fr/docs/irene.info.pdf> (you will need a TGCC login and password)
- Irene computing nodes: the nodes of partition **skylake** have 48 cores each.
  - Skylake nodes for regular computation (coming from irene.info)
  - Partition name: **skylake**
  - CPUs: 2x24-cores Intel Skylake@2.7GHz (AVX512)
  - Cores/Node: 48
  - Nodes: 1 656
  - Total cores: 79488
  - RAM/Node: 192GB
  - RAM/Core: 4GB
- 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.
- Compute nodes are diskless, meaning that `/tmp` is not hosted on a local hard drive anymore, but on system memory instead. It offers up to 16 GB (compared to 64 GB on Curie). Please note that any data written to it reduces the size of the memory that remains available for computations. For libIGCM post-treatment jobs you might need to use more cores than used at curie or, preferably, use `xlarge` nodes.
- 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. Since 2018/10/17 and libIGCM 1470, `xlarge` is proposed as default post-processing node.
  - 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

## 2. Job manager commands

- `ccc_msub job ->` submit a job
- `ccc_mdel 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.

## 3. Suggested environment

### 3.1. 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`. Do as follow:

```
cp ~igcmg/MachineEnvironment/irene/bashrc ~/.bashrc
cp ~igcmg/MachineEnvironment/irene/bashrc_irene ~/.bashrc_irene
```

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

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

```

module list
Currently Loaded Modulefiles:
 1) ccc                                10) mkl/17.0.6.256(default)          19) flavor/hdf5/parallel
 2) datadir/own(default)              11) flavor/buildcompiler/intel/17(default) 20) netcdf-c/4.6.0(default)
 3) dfldatadir/own(default)           12) intel/17.0.6.256(default)          21) netcdf-fortran/4.4.4(default)
 4) licshr/intel                      13) hwloc/1.11.3(default)             22) hdf5/1.8.20(default)
 5) c++/intel/17.0.6.256(default)      14) feature/openmpi/mpi_compiler/intel(default) 23) feature/bridge/heterogenous_mpm
 6) c/intel/17.0.6.256(default)        15) feature/openmpi/net/mxm(default)    24) nco/4.6.0(default)
 7) fortran/intel/17.0.6.256(default)  16) .tuning/openmpi/2.0(default)       25) cdo/1.7.2rc6(default)
 8) feature/mkl/lp64                  17) flavor/buildmpi/openmpi/2.0       26) ghostscript/9.19(default)
 9) feature/mkl/sequential             18) mpi/openmpi/2.0.2                 27) ferret/7.2(default)

```

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

### 3.2. 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
irene190 : 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/bashrc_irene`. 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.

## 4. File system

You have a main home where you arrive when connecting to irene, called "home de connexion" by the TGCC. You also have a home, a storedir, a workdir, a scratchdir by project. For example if you are working with project `gen2201` and `gen2212` you will have all following directories:

```

/ccc/cont003/home/~/login              # connexion home, where ***=your lab (lsce, ipsl, etc..)

/ccc/cont003/home/gen2201/login         # use it for sources, regular snapshot are in .snapshot
/ccc/cont003/home/gen2212/login

/ccc/store/cont003/gen2201/login
/ccc/store/cont003/gen2212/login

/ccc/work/cont003/gen2201/login
/ccc/work/cont003/gen2212/login

/ccc/scratch/cont003/gen2201/login
/ccc/scratch/cont003/gen2212/login

```

**IMPORTANT :** Check that you have read and write access to above directories (for your projects). Contact TGCC hotline if it is not the case.

**On the SCRATCH space any files that stays 60 days without being read or modified will be purged(deleted), as well as any directory that remains empty for 30 days.**

After connexion to irene, load your project environment as default using the module dfladatadir. For example if you will work on the project gen2201, do following (we suggest you to add the command into your .bashrc\_irene):

```
module switch dfladatadir dfladatadir/gen2201
```

By changing the dfladatadir, the variables \$CCCHOME, \$CCCWORKDIR, \$CCCSTOREDIR and \$CCCSCRATCHDIR point to the corresponding project directories. \$HOME is always the main connexion home.

You will also have new environment variables to access working directories :

```
GEN2201_ALL_CCCSCRATCHDIR=/ccc/scratch/cont003/gen2201/gen2201
GEN2201_CCCWORKDIR=/ccc/work/cont003/gen2201/login
GEN2201_ALL_HOME=/ccc/cont003/home/gen2201/gen2201
GEN2201_CCCSTOREDIR=/ccc/store/cont003/gen2201/login
GEN2201_CCCSCRATCHDIR=/ccc/scratch/cont003/gen2201/login
GEN2201_ALL_CCCWORKDIR=/ccc/work/cont003/gen2201/gen2201
GEN2201_HOME=/ccc/cont003/home/gen2201/login
GEN2201_ALL_CCCSTOREDIR=/ccc/store/cont003/gen2201/gen2201
```

If you previously worked at curie and your directories were in /cont003/dsm/login you will now find your data in a specific new project file system "dsmipsl". We recommend to move your data in your genci project file system. The TGCC hotline can help you if you want.

## 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 dfladatadir 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. A figure to understand file systems used on Irene

