

<http://autosubmit.readthedocs.io>

[Setup before the hands-on](#)  
[Tutorial](#)  
[Going further](#)

## Setup before the hands-on

You can follow the instructions below to setup the environment. After following these steps you'll be able to follow the tutorial.

### I. Virtual machine

#### A. Copy and import the virtual machine:

The virtual machine file is available online:

[ftp://autosubmit:sgH\\_123jHSS\\$Q@bscesftp.bsc.es/Autosubmit\\_EC\\_Earth\\_Hands\\_On.ova](ftp://autosubmit:sgH_123jHSS$Q@bscesftp.bsc.es/Autosubmit_EC_Earth_Hands_On.ova)

1. **Open it** (.ova file) with VirtualBox (File-->Import appliance) and check that everything is working. **The credentials are:**

User: autosubmit

Password: autosubmit

The image will also be provided during the session via usb pendrive, but it is strongly recommended to download it and prepare it before the session.

## II. MareNostrum3 users:

For the duration of this tutorial session you will be assigned with a user-ID and a password, which can be used to login to the MareNostrum III HPC at BSC.

1. **Replace** your user-ID in the `~/.ssh/config` in the following lines:

```
Host mn3
    HostName mn1.bsc.es
    User <mn3 user-ID>
    IdentityFile ~/.ssh/id_rsa
```

2. Now check you can login to MN3 without password. To do it, run the next command: `ssh mn3`

# Tutorial

## 1. First step: Configure Autosubmit in local host

Before creating any experiment, you need to configure your Autosubmit installation. There are two main paths that you will want to set up:

- Database path: Autosubmit uses a database to keep track of the experiments. This is the path that contains it.
- Experiment repository path: on this path Autosubmit will create the experiments' folders. This is the one that you will need to access to configure and monitor your experiments.

### a. Autosubmit environment configuration

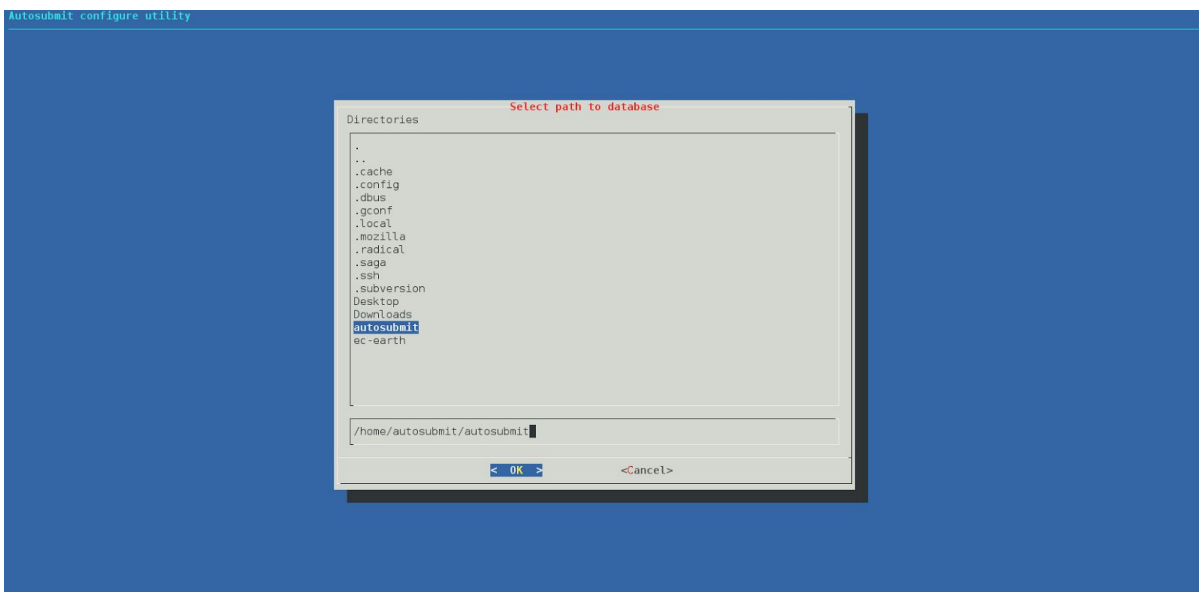
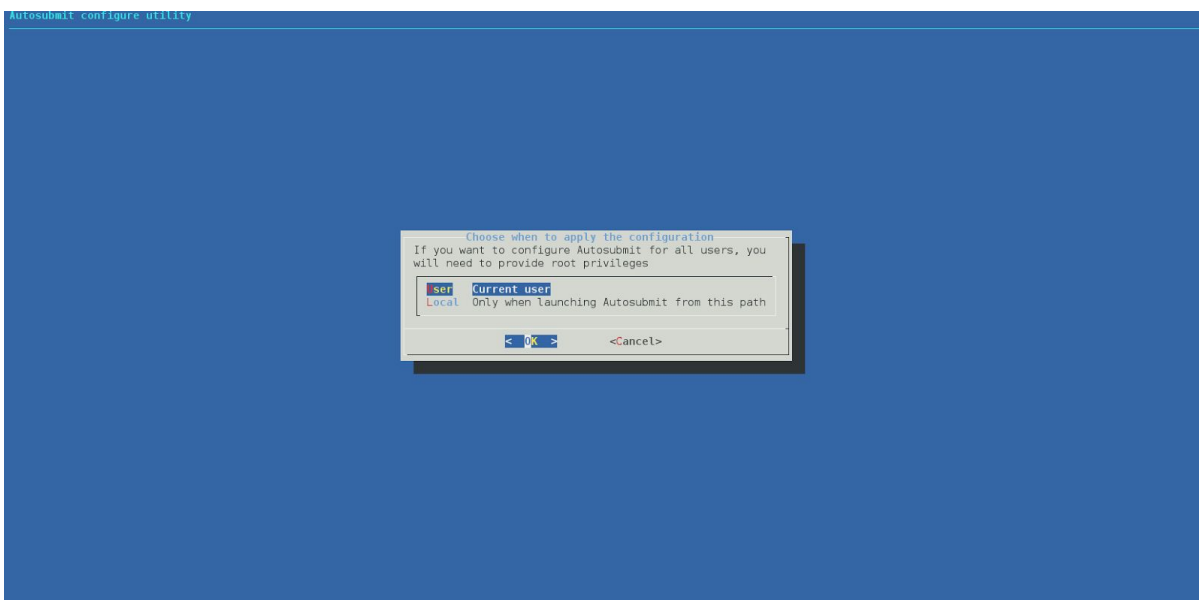
We will configure autosubmit to use a folder on your home directory. You will need to create it first:

```
mkdir ~/autosubmit
```

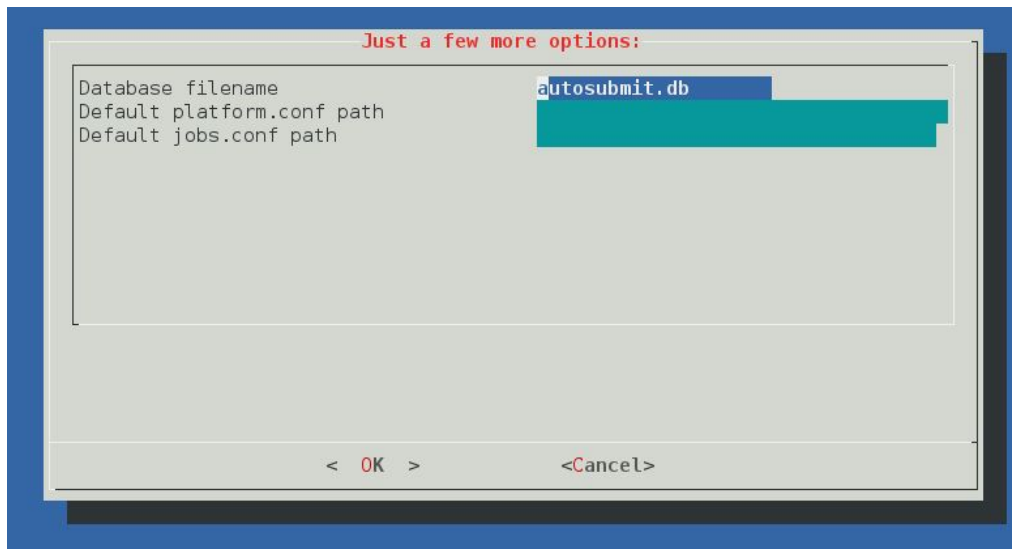
Next we will configure Autosubmit to use this folder to store the database and to use it as the experiments repository. The next command will configure it for your user:

### autosubmit configure

In the dialog window, choose option **'Current user'**, then select **'/home/autosubmit/autosubmit'** directory as a path to the database (and the same folder as the path to the experiments repository):



Keep the default values for the next steps. We will not use these options for this tutorial.



## b. Autosubmit environment installation

In this step we will create a new database for Autosubmit. Bear in mind that this command will overwrite your database if you have already created one, losing all the information it had. Be careful!

Run the command: `autosubmit install`

## 2. Second step: Create a dummy experiment

Now, you are going to create a new dummy experiment: an experiment that will run jobs using a template with just a `sleep()` call inside. It is a good way to test Autosubmit workflow capabilities and to test platform configurations.

### a. Autosubmit expid dummy

To register a new experiment you have to run the 'expid' command as below:

```
autosubmit expid -H marenostrum3 -d "Test for EC-Earth & Autosubmit training" -dm
```

Now you can check that in `~/autosubmit` there is a new folder named **'a000'** that will be the one that will store your experiment's configuration and monitoring files.

## b. Configure platform

The next step is to configure your platforms. On autosubmit you will always (by default) have a platform named LOCAL that corresponds to the machine running Autosubmit. For the HPC **you have to edit** the file:

`~/autosubmit/a000/conf/platforms_a000.conf` with your preferred text editor.

(You can keep all the commented lines (#) as a reference and copy & paste the configuration below these lines.)

You will have to use this configuration:

```
[marenostrum3]
TYPE = lsf
HOST = mn3
PROJECT = nct01
BUDGET = nct01:autosubmit_tutorial
USER = user-here
SCRATCH_DIR = /gpfs/scratch
```

From now on we will specify the paths relative to `~/autosubmit/a000`.

## c. Configure jobs

Now you need to configure the workflow. To do that, edit the file:  
**conf/jobs\_a000.conf**.

When you create a new experiment, Autosubmit defines some jobs as example. So, first of all, **remove all the jobs defined by default** on the jobs configuration file mentioned above.

(You can keep all the commented lines (#) as a reference and copy & paste the configuration below these lines.)

```
[LOCAL_SETUP]
FILE = copy-runtime.sh
PLATFORM = LOCAL

[SIM]
FILE = ece-ifs+nemo.sh
DEPENDENCIES = LOCAL_SETUP SIM-1
RUNNING = chunk
WALLCLOCK = 00:05
PROCESSORS = 2
```

#### d. Autosubmit create

Run `'autosubmit create a000'`. This command prepares the experiment to run.

If everything has worked well, you'll see something like this:

```

autosubmit@autosubmit:~/autosubmit/a000$ autosubmit create a000
Preparing .lock file to avoid multiple instances with same expid.

Checking configuration files...
autosubmit_a000.conf OK
platforms_a000.conf OK
jobs_a000.conf OK
expdef_a000.conf OK
Configuration files OK

Loading parameters...

Creating joblist...
Creating jobs...
Adding dependencies...
Removing redundant dependencies...


Saving joblist...

Plotting joblist...
Plotting...
Plot created at /home/autosubmit/autosubmit/a000/plot/a000_20160912_0846.pdf

Job list created succesfully
Remember to MODIFY the MODEL config files!

```

Then you will see the tree of your experiment's jobs in PDF format:



### e. Autosubmit run

Run 'autosubmit run a000'. Now the experiment will start:

### 3. Third step: Run EC-Earth:

At this step, you have tested your Autosubmit installation and your connection to Marenstrum3. The time to run EC-Earth with Autosubmit has arrived!!

#### a. Configure the experiment definition

Now you will transform your dummy experiment into one that will run two months of simulation with EC-Earth 3.2.0. Edit **conf/expdef\_a000.conf**:

```
DATELIST = 19900101
NUMCHUNKS = 2
CHUNKSIZE = 1
PROJECT_TYPE = local
PROJECT_PATH = /home/autosubmit/ec-earth/runtime/autosubmit
FILE_PROJECT_CONF = autosubmit.cfg
```

#### b. Configure the experiment jobs

On the previous run, we used a small wall clock and a small number of processors on the SIM jobs to reduce queuing time. Now you have to set up the real number of processors to use for your simulation. Edit again **conf/jobs\_a000.conf**:

```
[LOCAL_SETUP]
FILE = copy-runtime.sh
PLATFORM = LOCAL

[SIM]
FILE = ece-ifs+nemo.sh
DEPENDENCIES = LOCAL_SETUP SIM-1
RUNNING = chunk
WALLCLOCK = 00:30
PROCESSORS = 130
```



### c. Autosubmit create

Run `'autosubmit create a000'`. Autosubmit will create a copy of the runtime at the `'proj'` folder at this step and a new conf file that you will have to use.

### d. Configure proj\_xxxx.conf

Models can provide a custom configuration file for all parameters they need (e.g: `autosubmit.cfg`). This file will serve as a template for autosubmit to create the file `conf/proj_a000.conf`. Edit it:

```
NEM_NUMPROC = 64
IFS_NUMPROC = 64
```

### e. Run ec-conf

Go to the folder: `~/autosubmit/a000/proj/autosubmit`

Then check that the parameters below on `'marenostrum3'` section in your `'config-run.xml'` file are the following:

```
<Platform name="marenostrum3">
...
</Platform>

<Parameter name="ECEARTH_SRC_DIR">
  <Description>Base directory for EC-Earth sources</Description>
  <Type>PATH</Type>
  <Value>/gpfs/projects/nct00/nct00001/ec-earth/sources/sources</Value>
</Parameter>

<Parameter name="RUN_DIR">
  <Description>EC-Earth run directory</Description>
  <Type>PATH</Type>
  <Value>/gpfs/scratch/${GROUP}/${USER}/${exp_name}/${run_start_date}/${member
}/runtime</Value>
</Parameter>
```

```
<Parameter name="INI_DATA_DIR">  
  <Description>Initial data directory used by EC-Earth</Description>  
  <Type>PATH</Type>  
  <Value>/gpfs/projects/nct00/nct00001/ec-earth/inidata</Value>  
</Parameter>
```

Then just run ec-conf:

```
~/ec-earth/sources/util/ec-conf/ec-conf --platform  
marenostrum3 config-run.xml
```

This will create the appropriate EC-Earth run script depending on the values you put in the config-run.xml.

#### f. Autosubmit run

Execute `'autosubmit run a000'` and your experiment will start!

## Going further

Following exercises are optional. In the tutorial you have completed few steps to run a dummy experiment and an EC-Earth experiment using Autosubmit installation in a virtual machine. The simulations run in MareNostrum3 using temporal training accounts. If you want to try the tutorial using your HPC infrastructure and/or your own PC/Laptop, here below you have brief instructions to do so:

### A. If you are going to use your HPC:

- a. Configure your SSH access to the HPC on the `$HOME/.ssh/config`

e.g:

# contents of `$HOME/.ssh/config`

Host **hpc-name-here**

HostName **hpc-host-here**

User **hpc-user-here**

IdentityFile `~/.ssh/id_rsa`

- b. Please make also sure that your ssh keys are registered there
- c. Check that now you can login to the HPC by using 'ssh myhpc' without password
- d. On the step **2a**, run the command:

```
autosubmit expid -H your-hpc -d "Test for EC-Earth & Autosubmit training" -dm
```

- e. On the step **2b**, add a section for your HPC platform:

[your-hpc]

TYPE = your scheduler type

HOST = your host defined in `$HOME/.ssh/config`

PROJECT = your project name

BUDGET = by default, it is your project name. If that is the case, you can remove this line

USER = your user on the machine

SCRATCH\_DIR = machine's scratch folder path

- f. Follow similarly steps **2c**, **2d** and **2e** to complete a dummy experiment

**B. If you are going to run EC-Earth:**

- a. Compile EC-Earth 3.2.0 at the HPC you are going to use
- b. Follow similarly steps **in step 3** and in the step **3e**, edit the section for your HPC platform:

```
<Platform name="your-hpc">
    ...
</Platform>
```

Use your own paths for the EC-Earth sources and inidata.

Then run:

```
~/ec-earth/sources/util/ec-conf/ec-conf --platform your-hpc
config-run.xml
```

**C. If you are going to use your own PC/Laptop:**

- a. Install GraphViz
- b. Install dialog
- c. Install Autosubmit  
<http://autosubmit.readthedocs.io/en/latest/installation.html>
- d. Get a local copy of EC-Earth sources (in case you want to run EC-Earth, if you want to try dummy experiments it is not needed)
- e. Start the tutorial again.

**Get involved or contact us:**

Autosubmit GitLab:

<https://earth.bsc.es/gitlab/es/autosubmit>

Autosubmit Mailing List:

[autosubmit@bsc.es](mailto:autosubmit@bsc.es)