Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Autosumbit Basics Tutorial

Pablo Echevarria
Domingo Manubens

20th Sep 2018

What is Autosubmit?

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

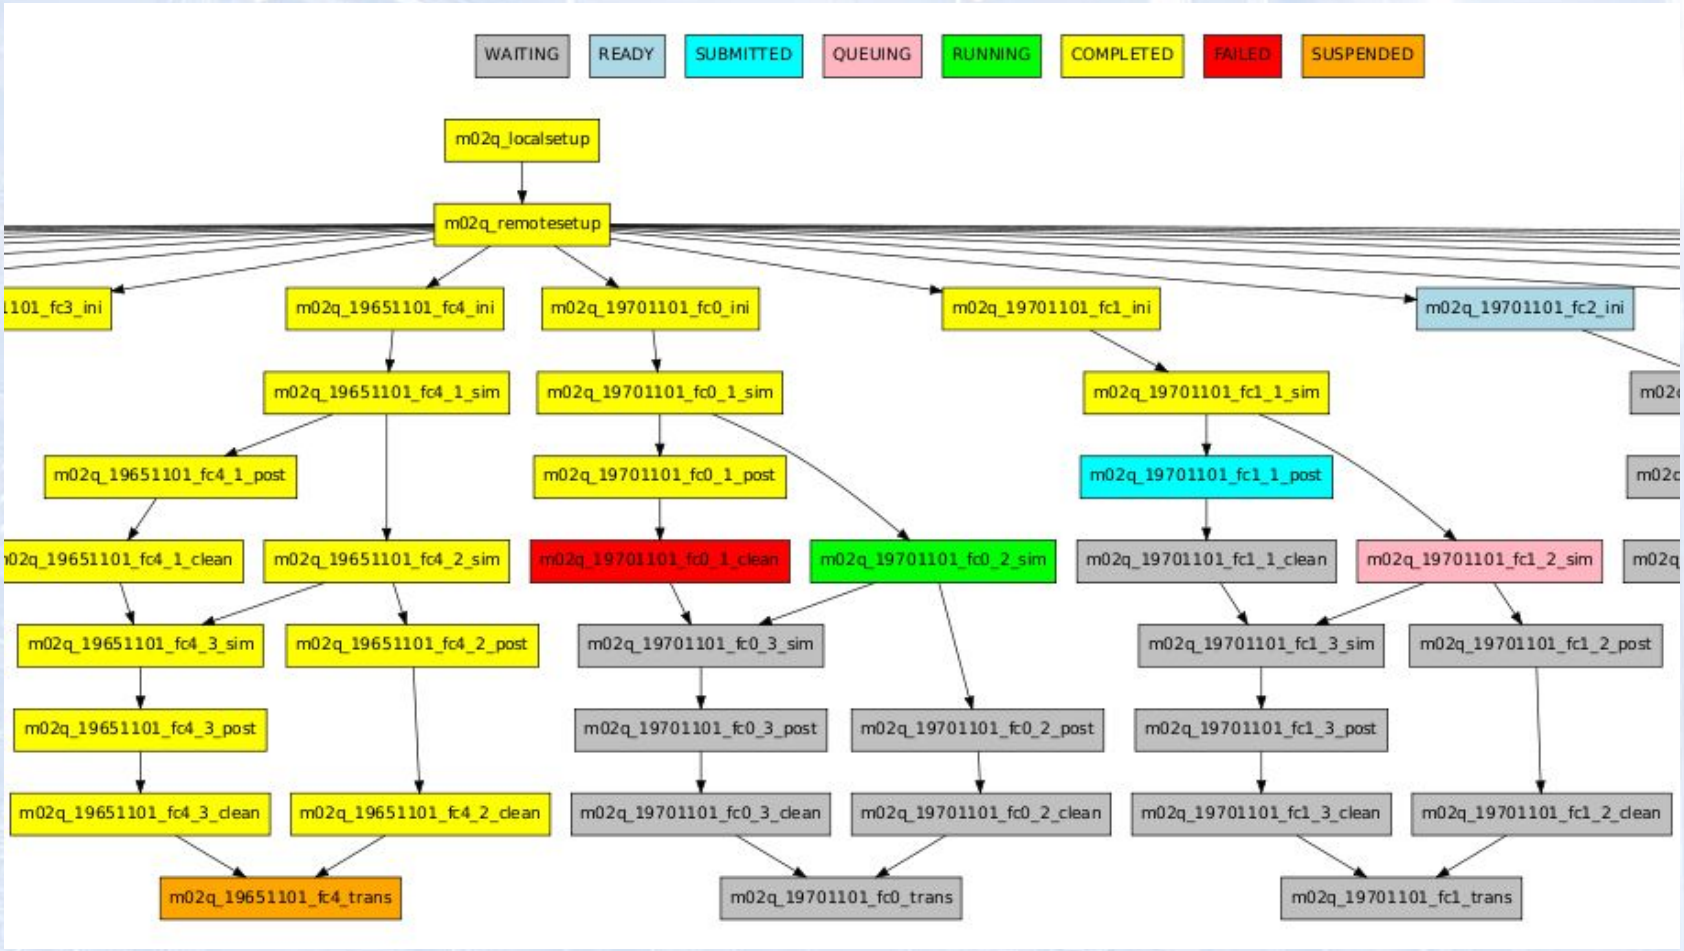
WAITING
READY
SUBMITTED
QUEUING
RUNNING
COMPLETED
FAILED
SUSPENDED
m02q_localsetup
m02q_remotesetup
m02q_19651101_fc4_ini
m02q_19701101_fc0_ini
m02q_19701101_fc1_ini
m02q_19701101_fc2_ini
m02q_19651101_fc4_1_sim
m02q_19701101_fc0_1_sim
m02q_19701101_fc1_1_sim
m02q_19651101_fc4_1_post
m02q_19701101_fc0_1_post
m02q_19701101_fc1_1_post
m02q_19651101_fc4_2_sim
m02q_19701101_fc0_1_clean
m02q_19701101_fc0_2_sim
m02q_19701101_fc1_1_clean
m02q_19701101_fc1_2_sim
m02q_19651101_fc4_3_sim
m02q_19651101_fc4_2_post
m02q_19701101_fc0_3_sim
m02q_19701101_fc1_3_sim
m02q_19701101_fc1_2_post
m02q_19651101_fc4_3_post
m02q_19701101_fc0_3_post
m02q_19701101_fc0_2_post
m02q_19701101_fc1_3_post
m02q_19651101_fc4_3_clean
m02q_19651101_fc4_2_clean
m02q_19701101_fc0_3_clean
m02q_19701101_fc0_2_clean
m02q_19701101_fc1_3_clean
m02q_19701101_fc1_2_clean
m02q_19651101_fc4_trans
m02q_19701101_fc0_trans
m02q_19701101_fc1_trans

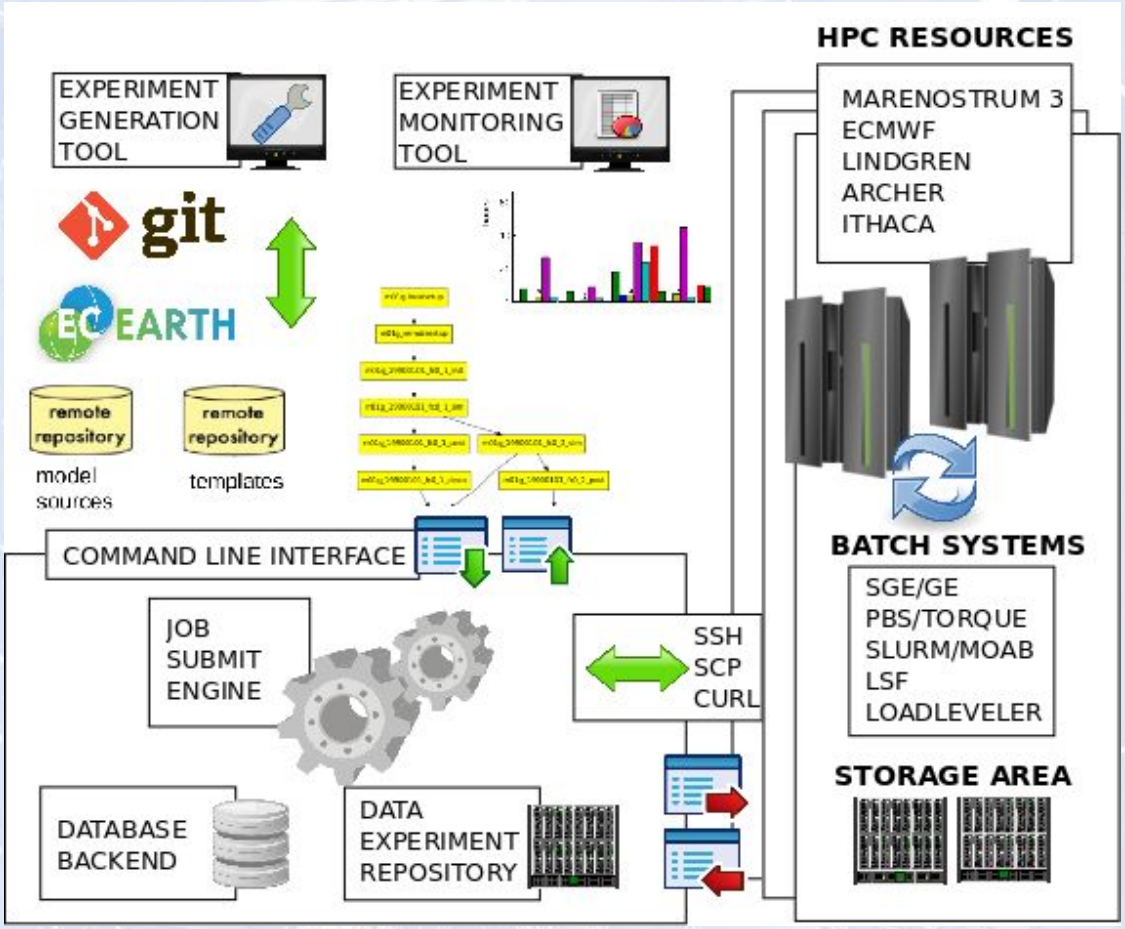

AUTOSUBMIT

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Why do you need a workflow manager?

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EC EARTH 3

1981

2015

1st Jan

1st Jan

10 years

4 members

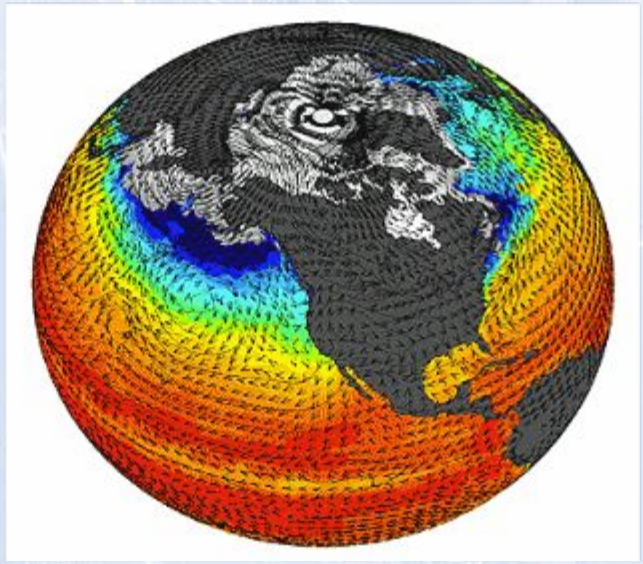

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Because when ...

- … you have been granted with computing resources on a new HPC and you need to deploy latest version of the model , …
- … you want to run long simulation for the 1850 to 2050 period, which lasts for 40 “human” days at HPC, …
- … and you want to interleave post processing tasks between each simulated year simulation, …
- … or you’re doing ensemble forecasting and need to add more members / startdates to an experiment, …
- … you detect an error in one of your chunks of data and need to recompute your job from the beginning, …

… you discover that …

- … you can’t remember all the scripts you need
- … you don’t remember the launch order of your scripts
- … you need the results soon and have to connect to launch jobs on weekends / nights

Advantages of a workflow manager

- Organization
- Monitoring
- Reproducibility
- Performance
- Efficient usage of resources

Part 1
Starting with Autosubmit

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The server(s)

We maintain Autosubmit stable version on Earth department workstations: https://earth.bsc.es/wiki/doku.php?id=computing:workstations
and 2 servers: https://earth.bsc.es/wiki/doku.php?id=library:bsc_servers

- You can create and monitor experiments using your workstation (it is faster to do it locally).

- Run experiments with **Autosubmit** server. Use bscesautosubmit01.bsc.es or bscesautosubmit02 for a more efficient data transfer to **/esarchive**

Warning !

- After simulations and Data Analysis, remove data files which are not needed anymore.
  - Storage costs money!

The dummy experiment

With Autosubmit you can create a dummy experiment to be used as the starting point of any other experiment we may develop.

You need to load the module

module load autosubmit

Check autosubmit version and help:

autosubmit -v

autosubmit -h

The dummy experiment

Register a dummy experiment

> **autosubmit expid --dummy --HPC LOCAL --description "Autosubmit basics tutorial running locally"**

The new experiment "a1do" has been registered.
Copying config files...
Experiment registered successfully
Remember to MODIFY the config files!

Go to **/esarchive/autosubmit/conf/a1do** and check files

platforms_**a1do**.conf
expdef_**a1do**.conf
jobs_**a1do**.conf

platforms_xxxx.conf

```
[marenostrum4]
# Queue type. Options: ps, SGE, LSF, SLURM, PBS, eceaccess
TYPE = slurm
HOST = mn1.bsc.es
PROJECT = bsc32
USER = bsc32704
SCRATCH_DIR = /gpfs/scratch
QUEUE = debug

[marenostrum4-dt]
# Queue type. Options: ps, SGE, LSF, SLURM, PBS, eceaccess
TYPE = ps
HOST = dt02.bsc.es
PROJECT = bsc32
USER = bsc32704
SCRATCH_DIR = /gpfs/scratch
```

jobs_xxxx.conf

s

```
[SYNC]
FILE = runtime/autosubmit/sync.sh
PLATFORM = marenostrum4-dt

[SIM]
FILE = runtime/autosubmit/ece-ifs+nemo.sh
DEPENDENCIES = SYNC SIM-1
RUNNING = chunk
WALLCLOCK = 00:20
PROCESSORS = 305
QUEUE = debug
```

expdef_xxxx.conf

```
[DEFAULT]
# Experiment identifier
# No need to change
EXPID = a1dp
# HPC name.
# No need to change
HPCARCH = marenostrum4

[experiment]
# Supply the list of start dates. Available formats: YYYYMMDD YYYYMMDDhh YYYYMMDDhhmm
# You can also use an abbreviated syntax for multiple dates with common parts: 200001[01 15] <=> 20000101 20000115
# 200001[01-04] <=> 20000101 20000102 20000103 20000104
# DATELIST = 19600101 19650101 19700101
# DATELIST = 1960[0101 0201 0301]
# DATELIST = 19[60-65]
DATELIST = 19900101
# Supply the list of members. Format fcX
# You can also use an abreviated syntax for multiple members: fc[0 1 2] <=> fc0 fc1 fc2
# fc[0-2] <=> fc0 fc1 fc2
# MEMBERS = fc0 fc1 fc2 fc3 fc4
# MEMBERS = fc[0-4]
MEMBERS = fc0
# Chunk size unit. STRING = hour, day, month, year
CHUNKSIZEUNIT = month
# Chunk size. NUMERIC = 4, 6, 12
CHUNKSIZE = 1
# Total number of chunks in experiment. NUMERIC = 30, 15, 10
NUMCHUNKS = 1
# Initial chunk of the experiment. Optional. DEFAULT = 1
CHUNKINI = 1
# Calendar used. LIST: standard, noleap
CALENDAR = standard
```

The dummy experiment

- Start and run the dummy experiment:

  > autosubmit create <your-expid-here>

  > autosubmit run

- Open a new terminal and obtain a graphical monitor plot:

  > autosubmit monitor <xxxx>

The dummy experiment

We can extend our experiment with 2 more members:

- Edit expdef_<xxxx>.conf and add fc1 and fc2
  - MEMBERS = fc0 fc1 fc2

- Re-create and recover dummy experiment:
  - autosubmit create <xxxx>
  - autosubmit recovery <xxxx> --all --save
  - autosubmit run <xxxx>

- Open a new terminal and obtain a graphical monitor plot:
  - autosubmit monitor <xxxx>

Check logs

- Autosubmit logs:
  - /esarchive/autosubmit/<expid>/tmp/
- Jobs logs:
  - /esarchive/autosubmit/<expid>/tmp/LOG_expid/

Part 2
Your first Autosubmit EC-Earth experiment on MareNostrum 4

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

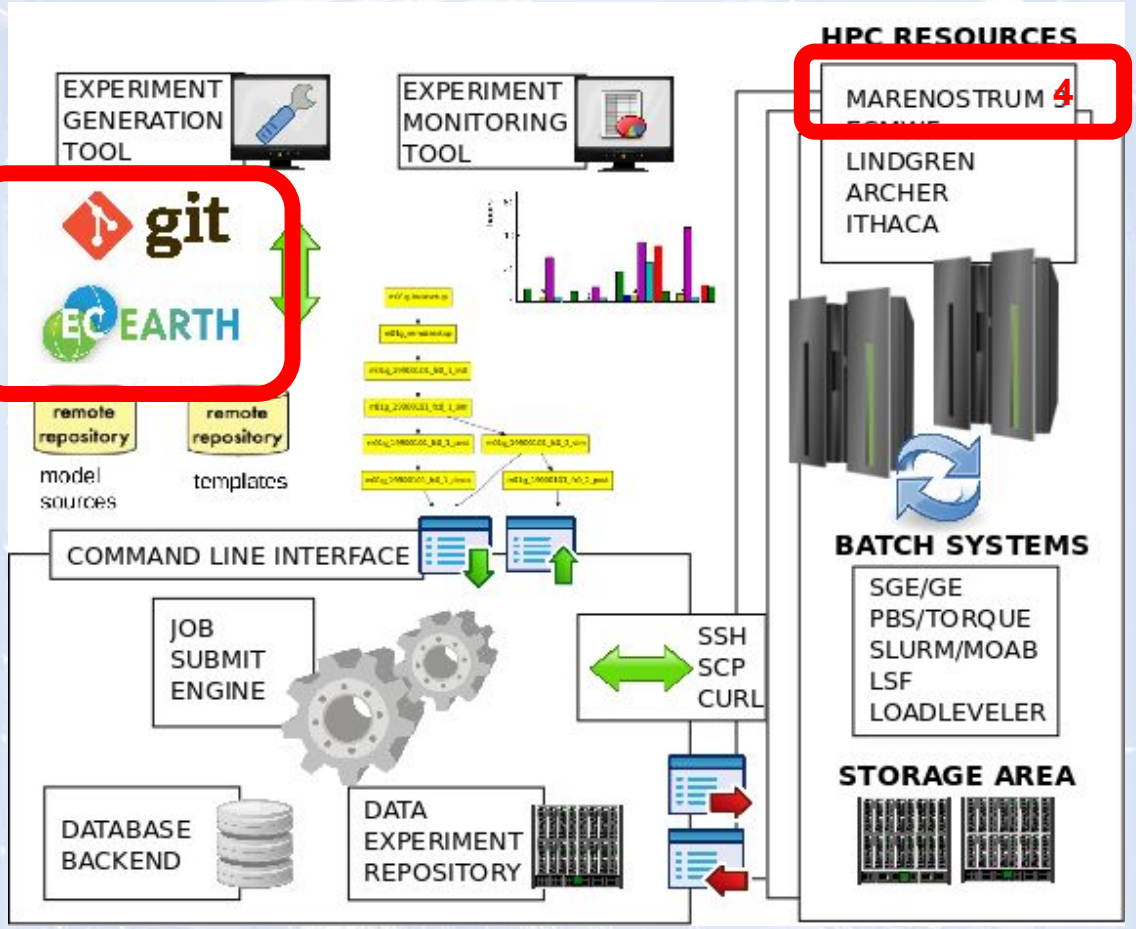
HPC RESOURCES
EXPERIMENT GENERATION TOOL
EXPERIMENT MONITORING TOOL
MARENOSTRUM 4
LINDGREN
ARCHER
ITHACA
git
EC EARTH
remote repository
remote repository
model sources
templates
COMMAND LINE INTERFACE
BATCH SYSTEMS
JOB SUBMIT ENGINE
SSH
SCP
CURL
SGE/GE
PBS/TORQUE
SLURM/MOAB
LSF
LOADLEVELER
STORAGE AREA
DATABASE BACKEND
DATA EXPERIMENT REPOSITORY

AUTOSUBMIT

EC-EARTH components

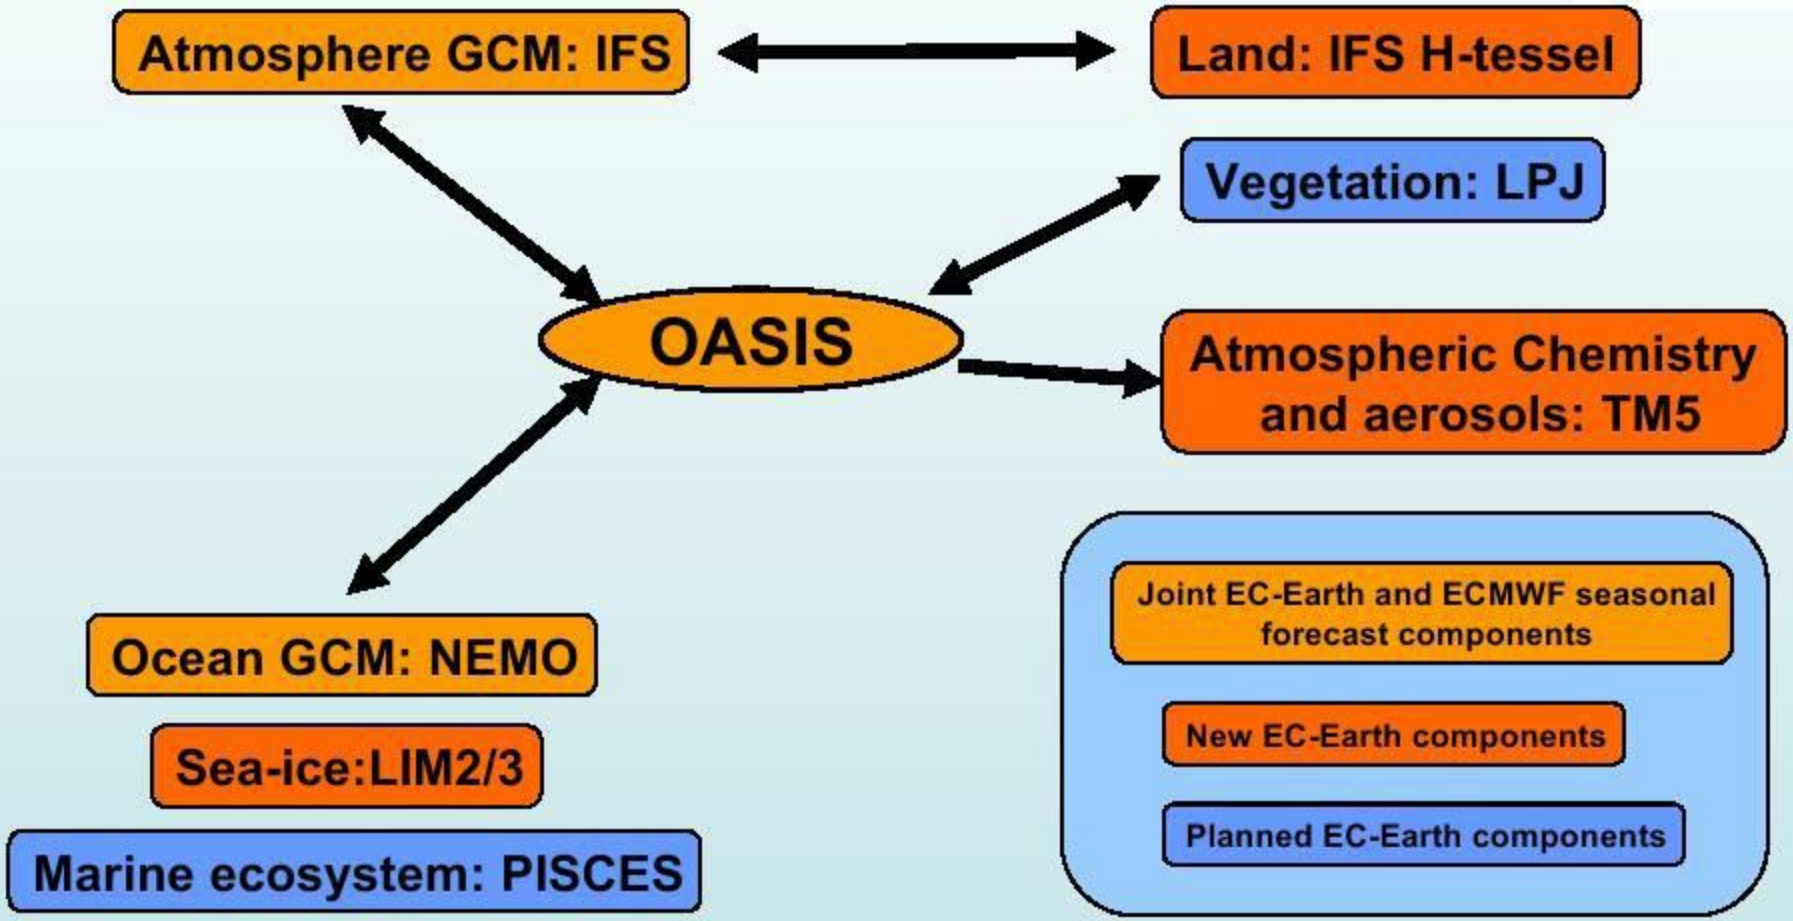

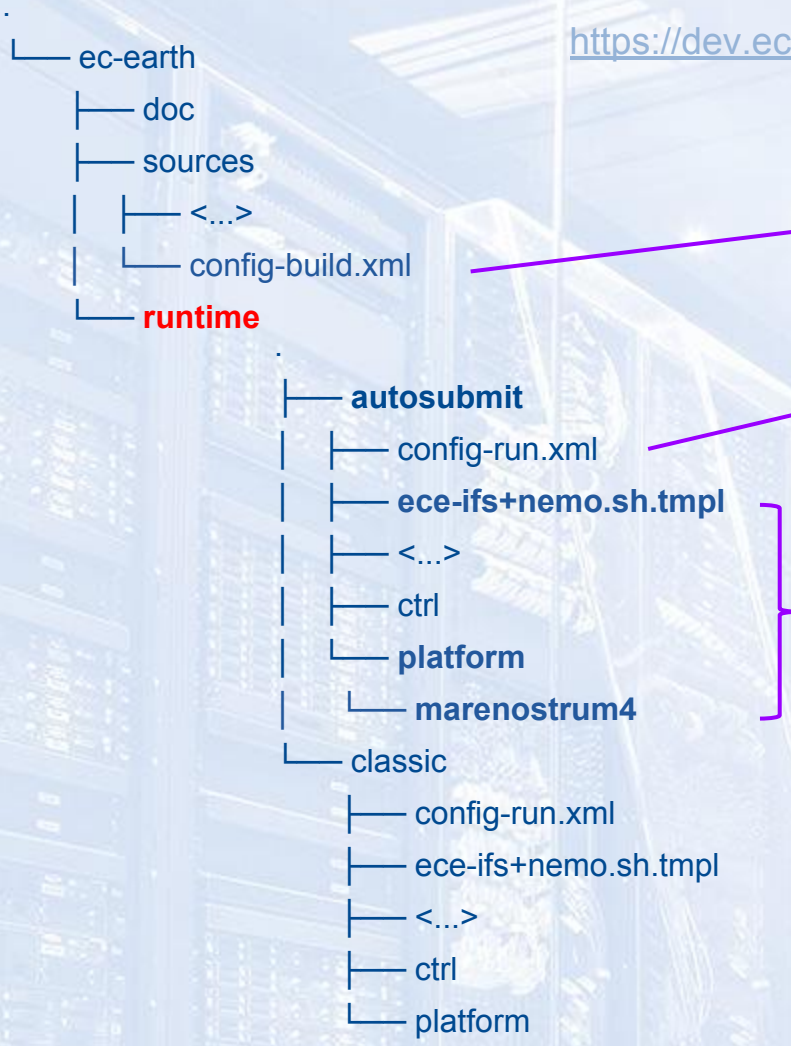

https://dev.ec-earth.org/projects/ecearth3/wiki/EC-Earth_3_User_Guide

Building
- Build Configuration
- Compiling

Running EC-Earth 3
- Run Configuration with ec-conf

Running EC-Earth 3 Experiments
- Running EC-Earth 3 on a specific platform

The EC-Earth experiment

Register a copy of a1dp experiment

> **autosubmit expid --copy a1dp --HPC marenostrum4 --description "Autosubmit basics tutorial running EC-Earth"**

The new experiment "a1do" has been registered.
Copying config files...
Experiment registered successfully
Remember to MODIFY the config files!

Go to **/esarchive/autosubmit/conf/a1do** and check files

platforms_**a1do**.conf → CHANGE TO YOUR MARENOSTRUM USER-ID !!!
expdef_**a1do**.conf
jobs_**a1do**.conf

Compare it with your dummy experiment

- Clone https://earth.bsc.es/gitlab/pechevar/utils.git

  Some tools to help you:

  propagate.py: iterate over the config files and compare using meld

  ./propagate.py --help to see how to use it

  example: call **module purge** before

  ./propagate.py a1dp a1dm

Other tools: **super_bashrc.sh, cca folder** and a lot of **garbage** please use under your responsibility, and please **report** any problem or possible improvement, it's started as a personal tools but could be turn a department toolkit

proj_xxxx.conf

```
[nemo]
# Number of parallel cores for OGCM component. NUMERIC = 16, 24, 36
NEM_NUMPROC = 128
# OGCM grid resolution. STRING = ORCA1L46, ORCA1L75, ORCA025L46, ORCA025L75 (NEMO)
NEMO_resolution = ORCA1L75
# Sea-Ice Model [Default: Do set "LIM2"]. STRING = LIM2, LIM3
ICE = LIM3

[ifs]
# Number of parallel cores for AGCM component. NUMERIC = 28, 100
IFS_NUMPROC = 128
# AGCM grid resolution, horizontal (truncation T) and vertical (levels L). STRING = T159L62, T255L62, T255L91, T511L91,
T799L62 (IFS)
IFS_resolution = T255L91

[xios]
# Number of parallel cores for IO servers. NUMERIC = 28, 100
XIO_NUMPROC = 48

[common]
# Model
MODEL = ecearth
# Version
VERSION = v3.2.3
```

The EC-Earth experiment

- Look at https://earth.bsc.es/gitlab/es/autosubmit_basic_runtime.git

- Start and run the experiment:

  > autosubmit create <your-expid-here>

  > autosubmit run

- Open a new terminal and obtain a graphical monitor plot:

  > autosubmit monitor <xxxx>

Checking your simulation

```
EXP01/
└── runtime
    ├── <...>
    ├── log
    │   └── 001
    ├── output
    │   ├── ifs
    │   │   └── 001
    │   └── nemo
    │       └── 001
    └── restart
        ├── ifs
        │   └── 002
        └── nemo
            └── 002
```

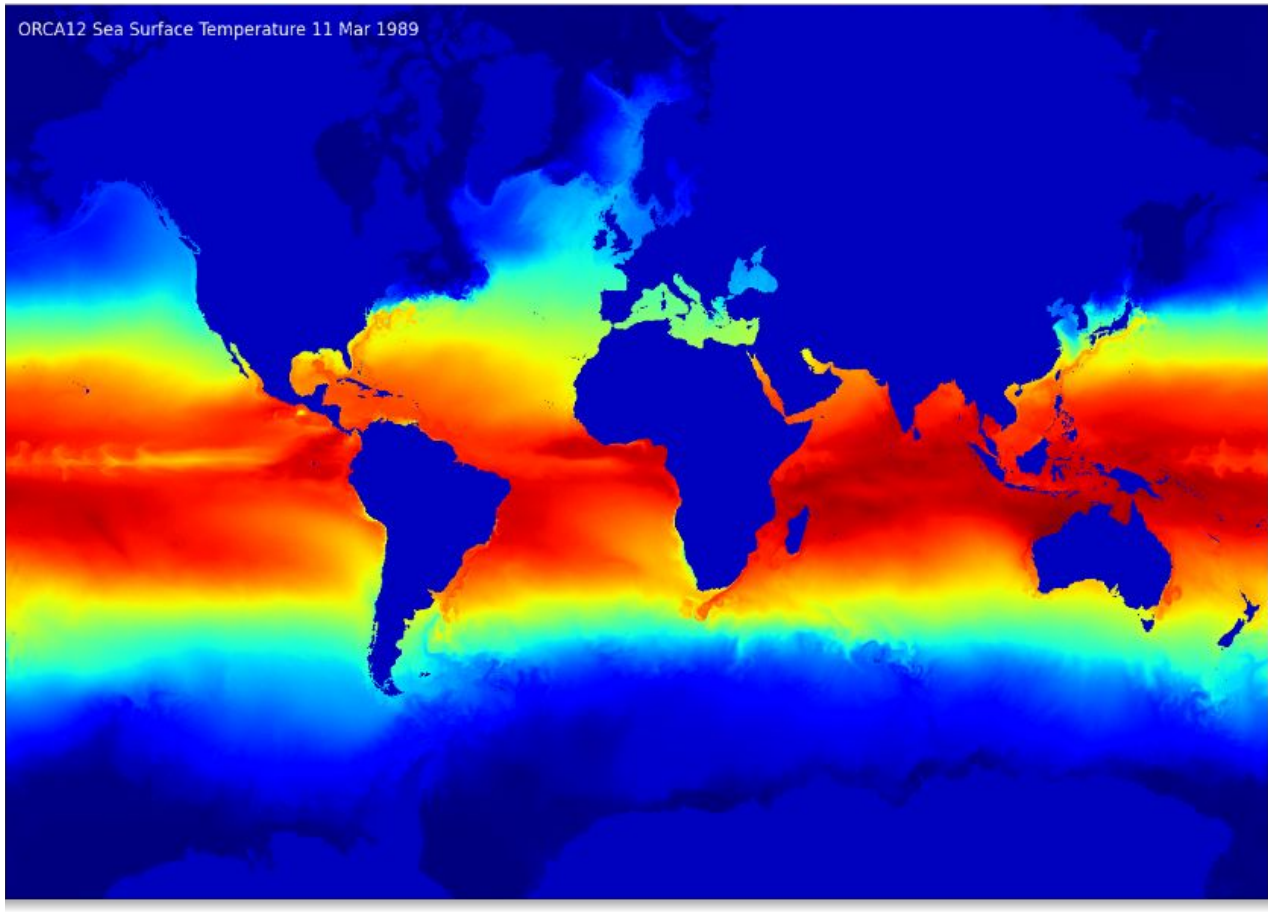

Autosubmit documentation

- Look at https://autosubmit.readthedocs.io/

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

domingo.manubens@bsc.es
pablo.echevarria@bsc.es