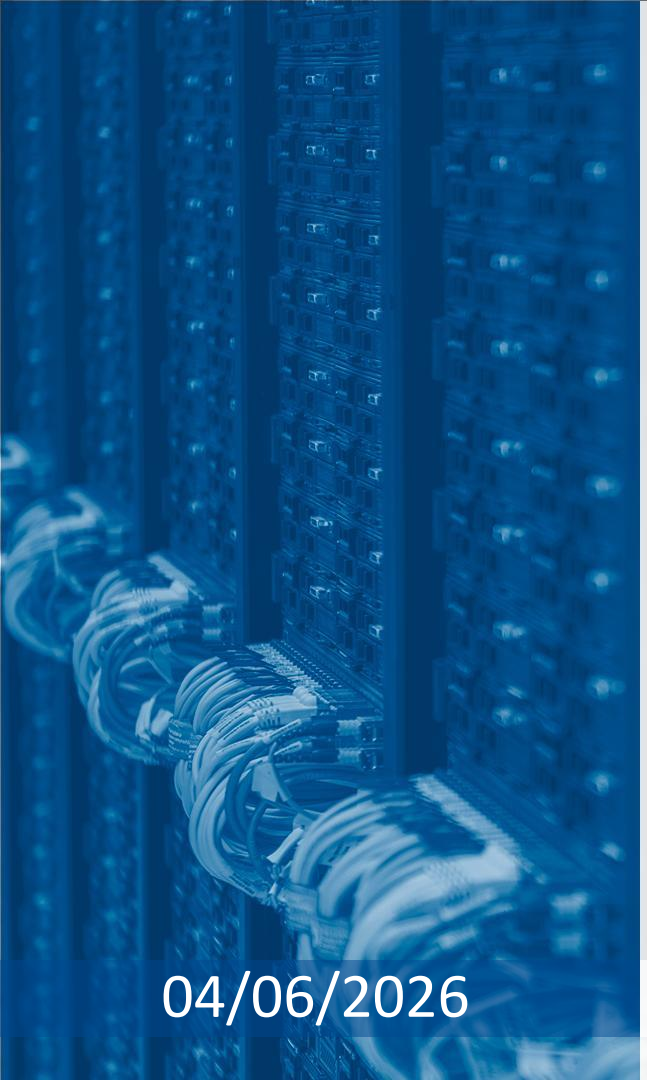

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

R user meeting

04/06/2026

Victòria A., Ariadna B.

Agenda

1. Ice-breaker: Persistent terminal sessions with ‘screen’
2. News
    - s2dv
    - CSTools
    - CSIndicators
    - CSDownscale
    - esviz
    - SUNSET
3. User presentation: Victòria Agudetse
4. Q&A

Ice-breaker: Persistent terminal sessions with ‘screen’

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Persistent terminal sessions with ‘screen’

1. Start an R job on a cluster
2. Close your laptop
3. Reconnect later and continue where you left off

```
screen -S myjob      # create session
screen -R myjob      # reconnect
screen -ls           # list sessions
```

- Keyboard shortcuts:

```
Ctrl+A D detach (session keeps running)
Ctrl+D exit (session ends if nothing else is running)
```

s2dv

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New release s2dv/2.3.0

Version 2.3.0 of ‘s2dv’ was released on 2026-04-30 and is now available on CRAN and installed on all BSC-ES machines.
The release includes:

- ★ New function DieboldMarianoTest() to compute the Diebold–Mariano test
- ★ New function TimescaleDecomposition() to decompose time series into temporal components
- ★ Several bugfixes and documentation improvement

issue: https://gitlab.earth.bsc.es/es/s2dv/-/issues/136
status: in production

CSTools

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New release CSTools/5.3.2

Version 5.3.2 of CSTools was released on 2026-05-28 and is now available on CRAN. It will be installed in BSC-ES machines in the coming weeks.
The release includes:

Developments:

★ CST_Subset(): New parameter `keep_metadata`.

Fixes:

★ CST_QuantileMapping(): Improve calls to functions from package 'qmap' to avoid problems in reverse dependencies

★ CST_QuantileMapping() no longer fails when a grid point has all-NA values

issue: https://gitlab.earth.bsc.es/external/cstools/-/issues/171
status: in production

CSIndicators

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New function: CST_SolarPVpot() and SolarPVpot()

- Computes solar photovoltaic (PV) potential as a capacity factor
- Expressed as the percentage of the maximum possible energy produced per PV grid cell
- Inputs:
  - `tas`: An s2dv_cube or array containing surface temperature.
  - `wind`: An optional s2dv_cube or array with surface wind speed (m/s).
  - `order`: A character string indicating the formula to calculate the PV potential indicator. If "second", 'wind' must be provided.
- Developed by Sara Moreno and Matías Olmo

MR: https://gitlab.earth.bsc.es/es/csindicators/-/merge_requests/94
status: in branch dev-SolarPVPot

CSDownscale

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New release CSDownscale/0.0.3

Version 0.0.3 of CSDownscale was released and submitted to CRAN on 2026-06-04. It will be installed in BSC-ES machines in the coming weeks.
The release includes:

- ★ `na.rm` parameter removed from Intbc(), as in the other CSDownscale functions. In CSDownscale, NA values are always removed before computation (`na.rm = TRUE`)
- ★ The default value of the `wet.day` parameter in Intbc() has been set to FALSE.
- ★ Errors caused by importing some functions from other packages have been resolved.
- ★ New citation in description file (Ramon et al, 2021).

issue: https://gitlab.earth.bsc.es/es/csdownscale/-/issues/16
status: in branch dev-0.0.3

esviz

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New release esviz/0.0.3

Version 0.0.3 of ‘esviz’ was released on 2026-05-08 and is now available on CRAN.

The release includes:

- ★ New function VizPolygonsMap() to plot polygon-based maps in different projections.
- ★ New parameter 'bkg_shapefile' in VizPolygonsMap() to include a secondary shapefile as background map.
- ★ Enable reading of shapefiles in VizEquiMap() and VizCombinedMap().
- ★ Added import of '%>%' from 'magrittr'.

issue: https://gitlab.earth.bsc.es/es/esviz/-/issues/37
status: in master

SUNSET

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

New indicator: FireIndex3030pXX

This indicator is similar to the fire index based on the **30-30-30 rule for wildfire risk** (temperature above 30ºC, RH below 30%, wind speed above 30 km/h), but uses a user-defined percentile pXX for the wind speed.

In the recipe, it should be requested as **"FireIndex3030pXX"**, replacing XX with the desired percentile. It can be computed with Indicators() or Crossval_indicators().

```
Indicators:
   name: FireIndex3030p66
   method: mean
   period: monthly
   start: '0601'
   end: '0831'
```

MR: https://gitlab.earth.bsc.es/es/sunset/-/merge_requests/301
status: in production

Fix redundant month labels

Fix `.get_plot_time_labels()` to avoid redundant labels when the start and end months of an aggregated time period are the same:

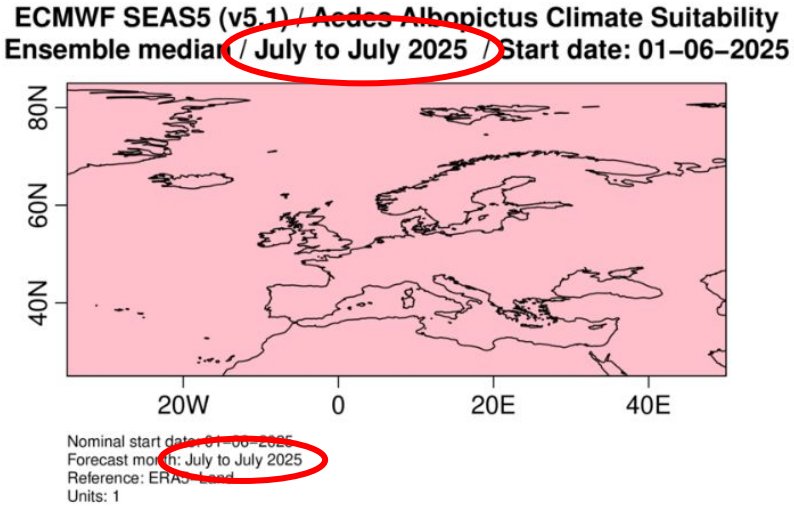

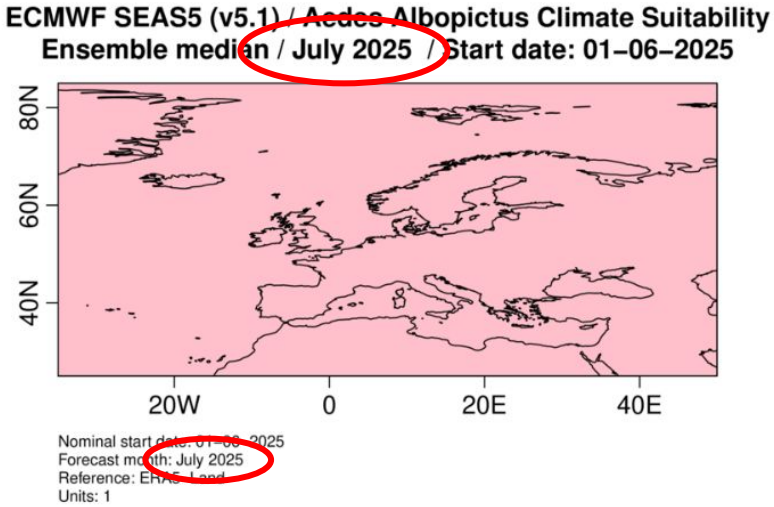

MR: https://gitlab.earth.bsc.es/es/sunset/-/merge_requests/327

status: in production

New caption developments

- Extra caption: costum text specified by the user in the recipe

```
Visualization:
  ...
  extra_caption: "Extra caption"
```

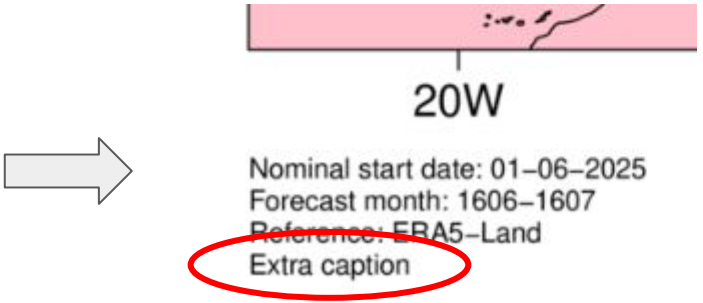

MR: https://gitlab.earth.bsc.es/es/sunset/-/merge_requests/307
status: in production

New caption developments

- Automatically include the indicator aggregation method (when applied)

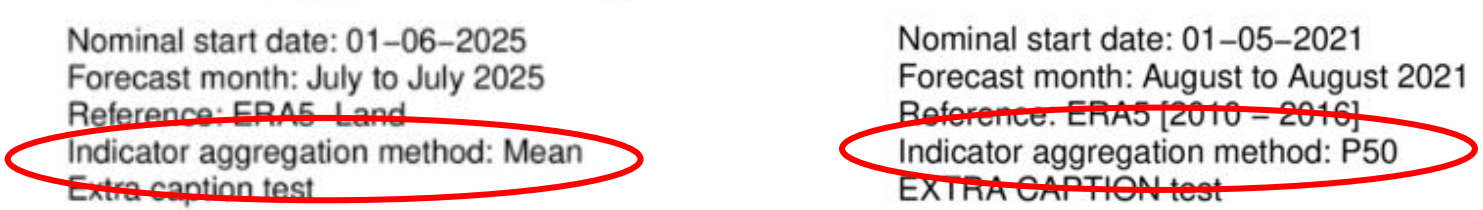

MR: https://gitlab.earth.bsc.es/es/sunset/-/merge_requests/307
status: in production

What do we know about BSC infrastructure?

Authors: Núria Pérez-Zanón and An-Chi Ho (December 2021)
Updated by Victòria Agudetse (June 2026)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC-ES Infrastructure: shared machines and tools

Aside from the data and software in our personal laptops, we all have access to a **shared BSC infrastructure**.

We access the BSC infrastructure:

- ★ When we connect to one of the **BSC-ES Hubs**
  - To ssh from windows: https://earth.bsc.es/wiki/doku.php?id=computing:sshwindows
  - To set up passwordless ssh connection: https://earth.bsc.es/wiki/doku.php?id=computing:sshkeyautologon
- ★ When we connect to one of the **servers** or **HPC machines** in BSC (MN5, Nord4, shiny server, etc.)
  - https://earth.bsc.es/wiki/doku.php?id=library:computing&s[]=computing

PARTITIONS

Each machine can access different files

When we connect to the BSC infrastructure, we find several **partitions**.
Each partition serves a different purpose and is accessible from different machines.

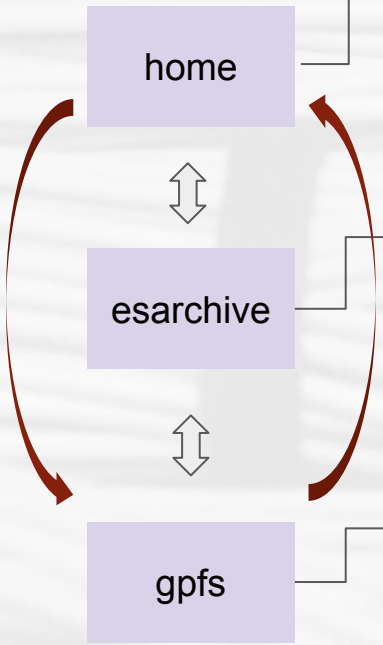

First entry point
/home/Earth/<username>/
Internet connection available
Daily back-up, limited space: 50 GB per user
Tips: https://earth.bsc.es/wiki/doku.php?id=computing:bsceshub
https://earth.bsc.es/wiki/doku.php?id=computing:sharedstorage

Data storage for BSC-ES department
Personal folder: /esarchive/scratch/<username>/
Internet connection available (from WS and Hub only)
NO back-up, limited total space for all users
Tip: add esarchive partition in a laptop (forbidden in Windows):
https://earth.bsc.es/wiki/doku.php?id=computing:addesarchive

General storage accessible from all HPC machines (Nord4, MN5)
Personal folder: /gpfs/scratch/bsc32/bsc032XXX/
No internet connection
Limited total space for all users in ES

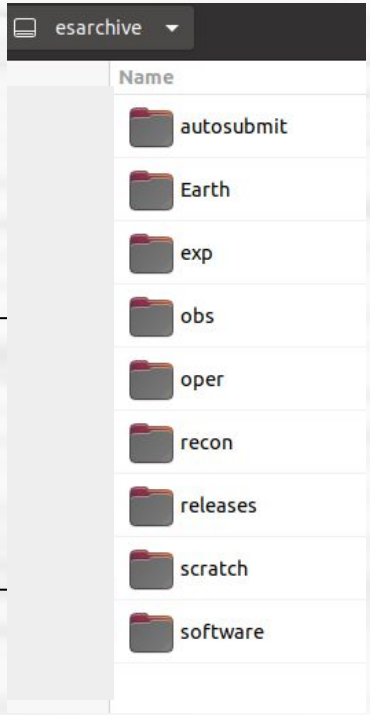

PARTITIONS

Other machines serve more specific purposes

It is also possible to connect to BSC infrastructure through **servers** (physical machines), which have different uses:

- **bscesdownloads01.bsc.es**
  - Download data
- **Transfer machine (transfer[1-4].bsc.es)**
  - Internal transfer of data, e.g. from esarchive to GPFS and vice versa.
- **bscesshiny01.bsc.es**
  - Shiny server, hosts shiny apps.
- **bscesftp.bsc.es**
  - Share files externally, see: https://earth.bsc.es/wiki/doku.php?id=computing:public_ftp
- **bscesautosubmit03.bsc.es and bscesautosubmit04.bsc.es**
  - Launch workflows with the Autosubmit workflow manager https://earth.bsc.es/wiki/doku.php?id=tools:autosubmit

SERVERS

BSC-ES has shared software stacks

A **software stack** is the collection of programs and modules (including the operating system, architectural layers, protocols, runtime environments, …) that are installed in a machine.

- ★ The software stack at BSC can be different between different machines and departments.
- ★ We have access to:
  - BSC software stack (not managed by CES)
  - BSC-ES software stack (managed by CES)
    - In some machines, we should edit the **bashrc** to use it (instructions are always in the wiki: https://earth.bsc.es/wiki/doku.php?id=library:computing)
    - It is built on **modules**, some useful commands are:
      - *module list* # show all loaded modules
      - *module load ** # load the '*' module
      - *module av ** # show all available modules matching '*'
- ★ Open an issue in **the Requests GitLab** to ask for new software or R packages

SOFTWARE STACK

What we need to know about each machine

- Access
- BSC-ES software stack
- Purpose
- Partitions
- Internet access
- Does it use a job scheduler? Which one?
- Available resources: Memory, nodes, cores per node…

Find the information here:
https://earth.bsc.es/wiki/doku.php?id=library:computing

BSC-ES Hubs

Summary

- ★ **Access:** `ssh <username>@bsceshub0X.bsc.es`
- ★ **Software stack:** module-based.
  For R, use: R-bundle-CRAN/2024.06-foss-2023b.
- ★ **Purpose:** Machines for **light tasks**, like debugging code (small data) or running startR workflows in remote machines.
- ★ **Partitions:** home and esarchive
- ★ **Job scheduler:** None
- ★ **Internet connection:** Yes

More info: https://earth.bsc.es/wiki/doku.php?id=computing:bsceshub

MACHINES

Nord4

Summary

- ★ **Access:** `ssh bscXXXXX@n4loginX.bsc.es`
- ★ **Software stack:** module-based.
  For R, use: R-bundle-CRAN/2023.12-foss-2020b
- ★ **Purpose:** High Performance Computing (HPC) cluster for **memory-intensive jobs**.
- ★ **Partitions:** esarchive and gpfs
- ★ **Job scheduler:** Slurm
- ★ **Internet connection:** Only on n4login0. Not available when running jobs.

More info: https://earth.bsc.es/wiki/doku.php?id=computing:nord4

MareNostrum5

Summary

- ★ **Access:** `ssh bscXXXXX@gloginX.bsc.es` or `ssh bscXXXXX@aloginX.bsc.es`
- ★ **Software stack:** modules and conda environments.
  For R, use: R-bundle-CRAN/2023.12-foss-2023b
- ★ **Purpose:** High Performance Computing (HPC) cluster for **memory-intensive jobs** and jobs requiring GPUs (massively parallel computing).
- ★ **Partitions:** gpfs
- ★ **Job scheduler:** Slurm
- ★ **Internet connection:** Only on glogin4 and alogin4. Not available when running jobs.

More info: https://earth.bsc.es/wiki/doku.php?id=computing:mn5

MACHINES

What is Slurm and how to use it

Slurm is a **job scheduler** that distributes computing resources between the HPC machine users.

★ We can launch a script <job_script> with the following command:

```
sbatch <job_script> -A <account> -q <qos> [-t <hh:mm:ss>] [-n <cpu_number>] [other arguments]
```

★ Or request a session to work interactively:

```
salloc -A <account> -q <qos> [-t <hh:mm:ss>] [-n <cpu_number>] [other arguments]
```

MACHINES

Slurm mandatory arguments

```
sbatch <job_script> -A <account> -q <qos> [-t <hh:mm:ss>] [-n <cpu_number>] [other arguments]
```

Because Nord4 and MN5 are accessible outside BSC-ES (and even BSC), the account and qos are mandatory arguments we always need to provide:

Argument name	Meaning	Nord4 Options	MN5 Options
-A/--account	Account group	bsc32	bsc32
-q/--qos	Quality of service	bsc_es debug interactive	gp_bsces gp_debug gp_interactive acc_bsces acc_debug acc_interactive

Slurm mandatory arguments

```
sbatch <job_script> -A <account> -q <qos> [-t <hh:mm:ss>] [-n <cpu_number>] [other arguments]
```

Other optional arguments allow us to specify which resources we want for our job:

Argument name	Meaning	Nord4 Options	MN5 Options
-t/--time	Time limit	hh:mm:ss	hh:mm:ss
-n/--ntasks	Number of processors	Up to 48 per node	Up to 112 per node (general purpose)
--constraint=<constraint>	Request for special type of node	highmem	highmem
--exclusive	Do not share my node with other jobs	--exclusive	--exclusive

MACHINES

Where to find HPC machine information

- ★ For technical details read the **User Guides**:
  - MareNostrum 5 | BSC Support Knowledge Center
  - Nord4 | BSC Support Knowledge Center.
- ★ Monitor the status of your jobs with the **squeue** command. You can use flags such as squeue --start or squeue --long to see more information.
- ★ For real-time information use the **HPC User Portal:**
  - Check job stats: HPC User Portal - Home
  - Check machine status: HPC User Portal - Machines Stats

Asking for less will make your jobs run faster

- ★ Use the **debug queues** for smaller and quicker jobs.
- ★ More time limit + more cores = **longer queueing time**.
- ★ Special requests like --constraint=highmem or --exclusive also mean longer queueing time
- ★ The **RAM available for the job is proportional to the number of cores requested**; around ~2GB per core. Asking for half of the available cores will get you half of the total RAM in a node.

What do we know about BSC-ES infrastructure?

RECOMMENDATIONS

Recommendations

★ **Use git and save your scripts in GitLab (intermediate and final versions)**
- In an existing GitLab project or in personal project (create one here: https://gitlab.earth.bsc.es/projects/new)
- Documentation: https://earth.bsc.es/wiki/doku.php?id=library:computing#git
- If you have internet connection, you can source code directly from GitLab
- Clone repositories under /esarchive/scratch/<username>/

★ **Don't install local versions of R packages**
- If you do, we cannot debug the code and reproduce the errors
- Better to open an issue in Requests: it's easier to debug and everyone can use it

★ **If you find outdated information, let CES know!**

What do we know about BSC-ES infrastructure?

Q&A: What else do we need to know? What questions do we have?

- **Q: Where can we find information about the BSC-ES packages and tools?**
  A: BSC-ES in-house tools are listed in https://earth.bsc.es/wiki/doku.php?id=tools:tools&s[]=tools
  In-house R packages are listed in https://earth.bsc.es/wiki/doku.php?id=tools:rtools&s[]=tools and their documentation can be found on CRAN (search "CRAN <package name>") or their corresponding GitLab pages.
- **Q: How can we know what external tools are available?**
  A: You can check for availability of a specific software with module av <name> or you can type the name in the Requests GitLab search bar to see if someone else requested it before.

Thanks for joining

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Next meeting: TBA