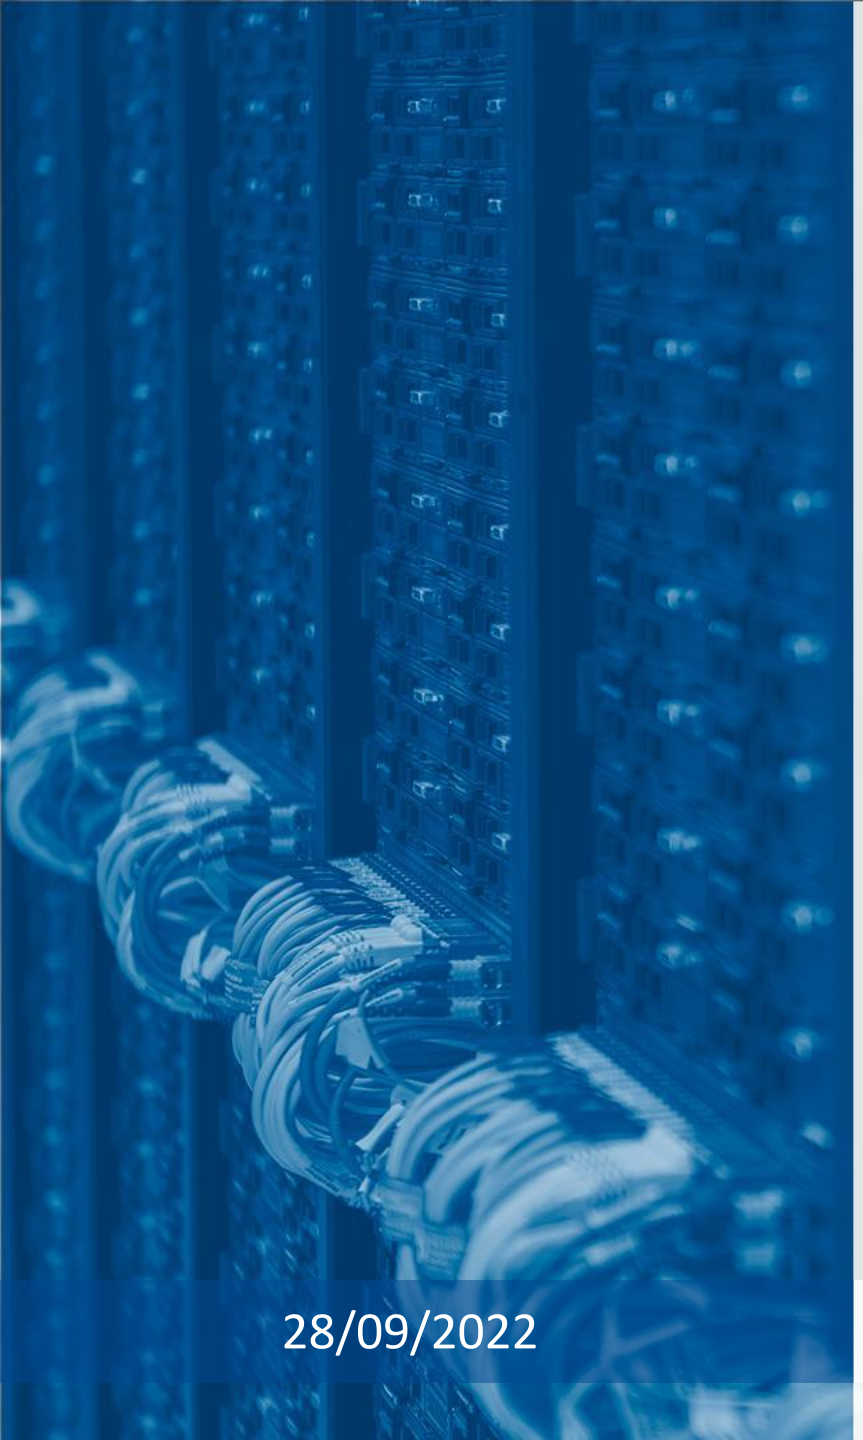

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Extending XIOS lossy compression functionalities using SZ

Xavier Yepes-Arbós, BSC
Sheng Di, ANL
Kim Serradell, BSC
Franck Cappello, ANL
Mario C. Acosta, BSC

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

The research leading to these results has received funding from the EU H2020 Framework Programme under grant agreement no. 823988
This material reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains

28/09/2022

14th JLESC Workshop, Urbana, IL, USA

Introduction

- The increase in the forecast capability of Earth System Models (ESMs) is strongly linked to the spatial resolution to solve more complex problems.
- This requires a large demand of computing power and it might generate a massive volume of model output which implies:
  - Data must be **efficiently** written into the storage system.
  - No more **offline post-processing** is affordable due to the size of the “raw” data.
  - A **high cost** of storage systems due to the huge data size.
- In this new exascale era, it will be mandatory to implement an efficient I/O management for ESMs.

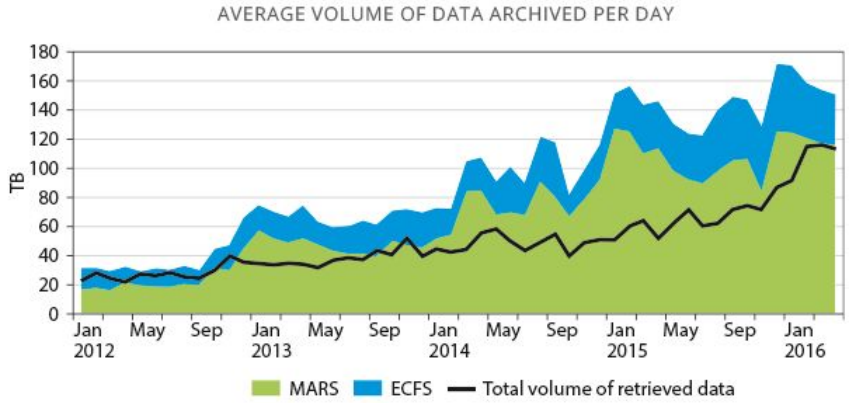

Figure source: ECMWF

The XIOS I/O server

- The I/O issue is typically addressed by adopting scalable parallel I/O solutions.
- In the climate community, a widely I/O tool used is **XIOS**.
- The XML Input/Output Server (XIOS) is an asynchronous MPI parallel I/O server developed by the Institut Pierre-Simon Laplace (IPSL).
- XIOS has the following features needed for climate modelling:
  - Output files are in **netCDF** format.
  - Written data is **CMIP-compliant** (CMORized).
  - It is able to post-process data **inline** to generate **diagnostics**.
- XIOS offers **lossless data compression** using gzip through HDF5.

Test case: OpenIFS and XIOS integration

- **OpenIFS** is an atmospheric general circulation model developed and maintained by the European Centre for Medium-Range Weather Forecasts (ECMWF).
- In the past we integrated XIOS into OpenIFS to address the former inefficient sequential I/O scheme (Yepes-Arbós et al.,2022).
- Although the **overhead** of outputting data through XIOS is really **small** for current resolutions, in the **future** this may well become a **bottleneck** because of the exponential growth of the output volume.
- The default lossless compression filter of HDF5 **does not fit** our needs:
  - If compression ratio is high, it takes too much **time**.
  - If it takes a reasonable amount of time, **compression ratio** is not enough.

OpenIFS-XIOS integration scheme

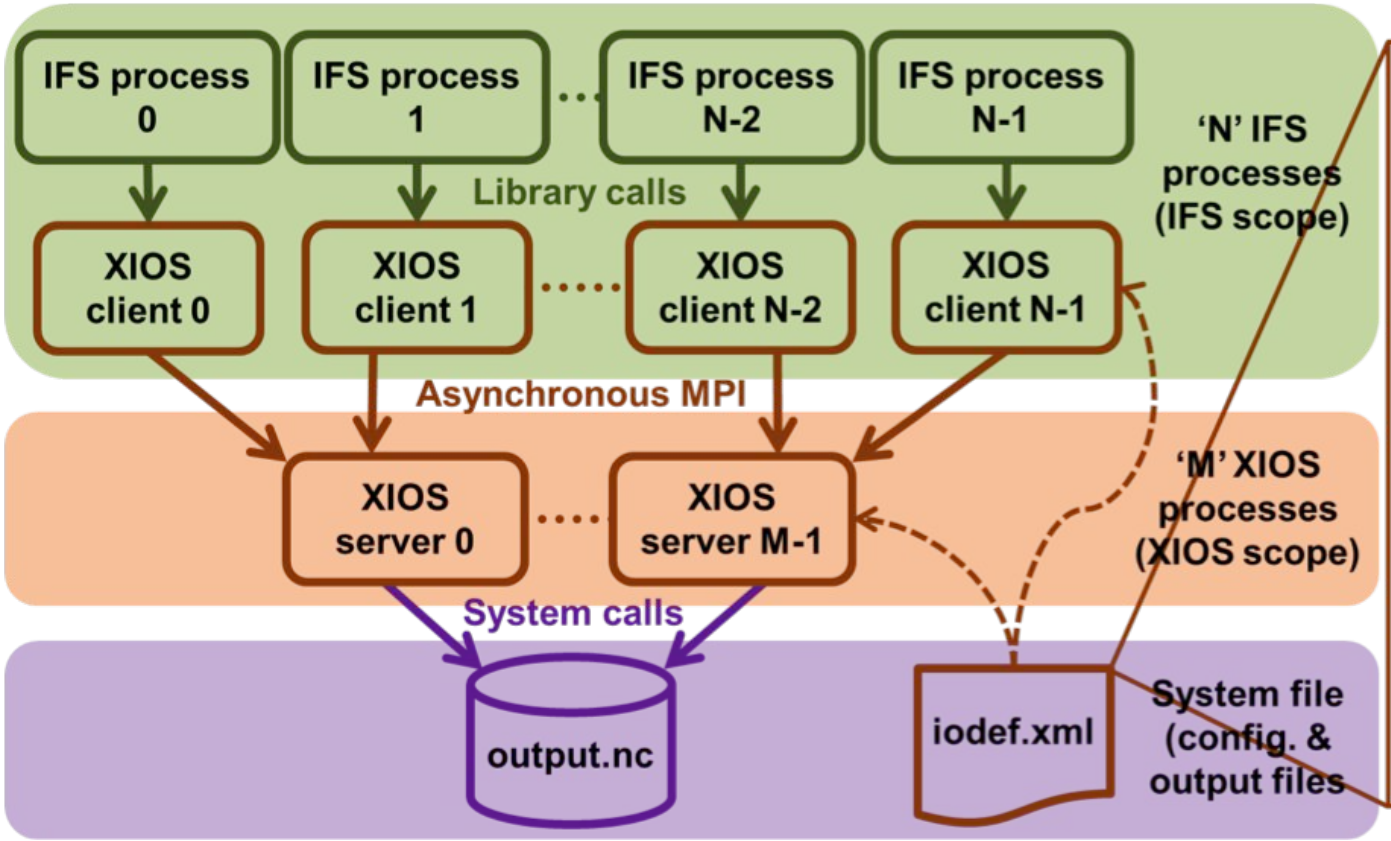


```
<grid id="grid_3d">
   <domain id="domain_2d"
   type="rectilinear" ni_glo=… />
   <axis id="axis_1d" n_glo=… />
</grid>

<field_definition>
   <field id="temp" grid_ref="grid_3d" />
   <field id="pressure" grid_ref="grid_3d" />
</field_definition>

<file_definition>
   <file name="output" freq_output="1mo">
      <field field_ref="temp" />
      <field field_ref="pressure" />
   </file>
</file_definition>
```

SZ lossy compression filter in XIOS

- We successfully tested the **SZ lossy compressor** from the Argonne National Laboratory (ANL) in XIOS with these advantages:
  - Reach high **compression ratios**.
  - Enough compression **speed** to considerably mitigate the I/O overhead.
  - Keep high **accuracy**.
- The SZ compressor is registered as a third-party filter of HDF5 which facilitates the integration in XIOS:

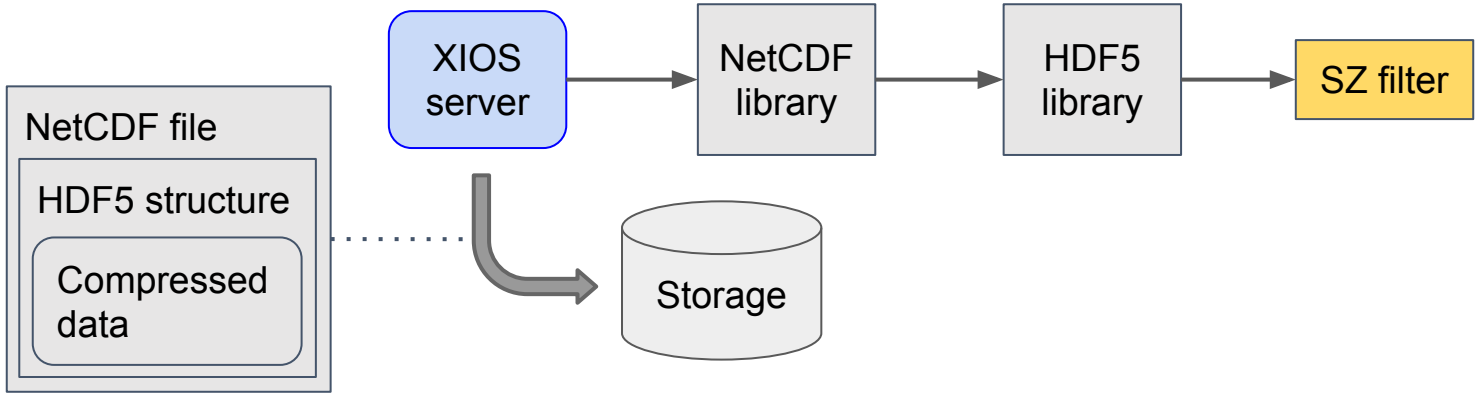

- However, we needed and wanted to implement other features.

Added compression functionalities in XIOS

- In climate modelling accuracy varies between fields.
- ANL added support to set different compression parameters per field.
- XIOS source code has been modified to support **independent compression parameters** per field through the XML configuration files:

```
<file id="file3" name="AMIP_6h_reduced_ml" output_freq="6h" output_level="10" compression_type="lossy" error_bound_mode="abs" error_bound="0.001" enabled="true" >
  <field field_ref="pres" name="pres" freq_op="6h" operation="instant" />
  <field field_ref="t" name="t" freq_op="6h" operation="instant" error_bound="0.1" />
  <field field_ref="u" name="u" freq_op="6h" operation="instant" error_bound_mode="rel" error_bound="0.01" />
  <field field_ref="v" name="v" freq_op="6h" operation="instant" error_bound_mode="rel" error_bound="0.01" />
  <field field_ref="q" name="q" freq_op="6h" operation="instant" error_bound="0.00001" />
  <field field_ref="clwc" name="clwc" freq_op="6h" operation="instant" compression_type="lossless" compression_level="4" />
  <field field_ref="ciwc" name="ciwc" freq_op="6h" operation="instant" compression_type="lossless" compression_level="4" />
  <field field_ref="cc" name="cc" freq_op="6h" operation="instant" compression_type="lossless" compression_level="6" />
</file>
```

Added compression functionalities in XIOS (2)

According to the command “ncdump -h -s AMIP_6h_reduced_ml.nc”:

- pres:_Filter = "32017,2,0,271360,91,0,981668463,0,0,0" ;
- t:_Filter = "32017,2,0,271360,91,0,1036831949,0,0,0" ;
- u:_Filter = "32017,2,0,271360,91,1,0,1008981770,0,0" ;
- v:_Filter = "32017,2,0,271360,91,1,0,1008981770,0,0" ;
- q:_Filter = "32017,2,0,271360,91,0,925353388,0,0,0" ;
- clwc:_DeflateLevel = 4 ;
- ciwc:_DeflateLevel = 4 ;
- cc:_DeflateLevel = 6 ;

```
netcdf AMIP_6h_reduced_ml {
dimensions:
        axis_nbounds = 2 ;
        cell = 271360 ;
        nvertex = 4 ;
        model_levels = 91 ;
        time_counter = UNLIMITED ; // (4 currently)
```

Added compression functionalities in XIOS (3)

- Since HDF5 1.10.2 there is support for **compression using parallel I/O**.
- We have adapted the XIOS source code to take advantage of this, both for lossless and lossy compression.
- Although it works, the performance is not very good. Performance improvements introduced in experimental release 1.13.1.

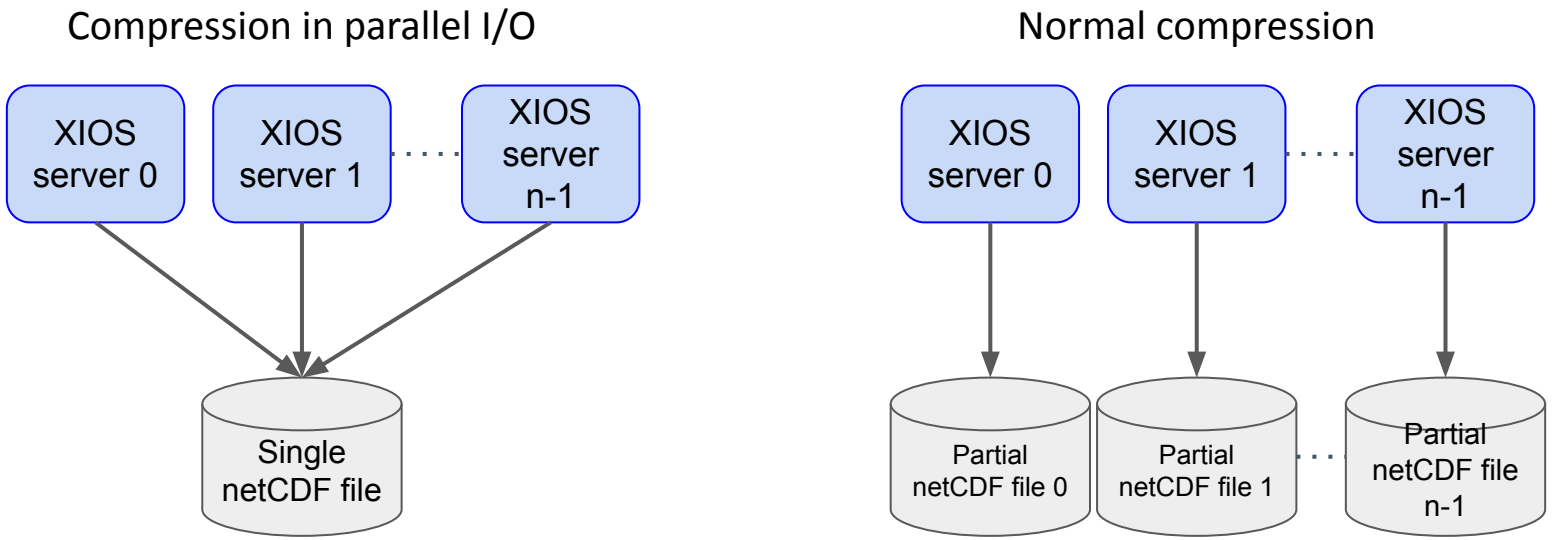

gzip vs. SZ compression: preliminary results

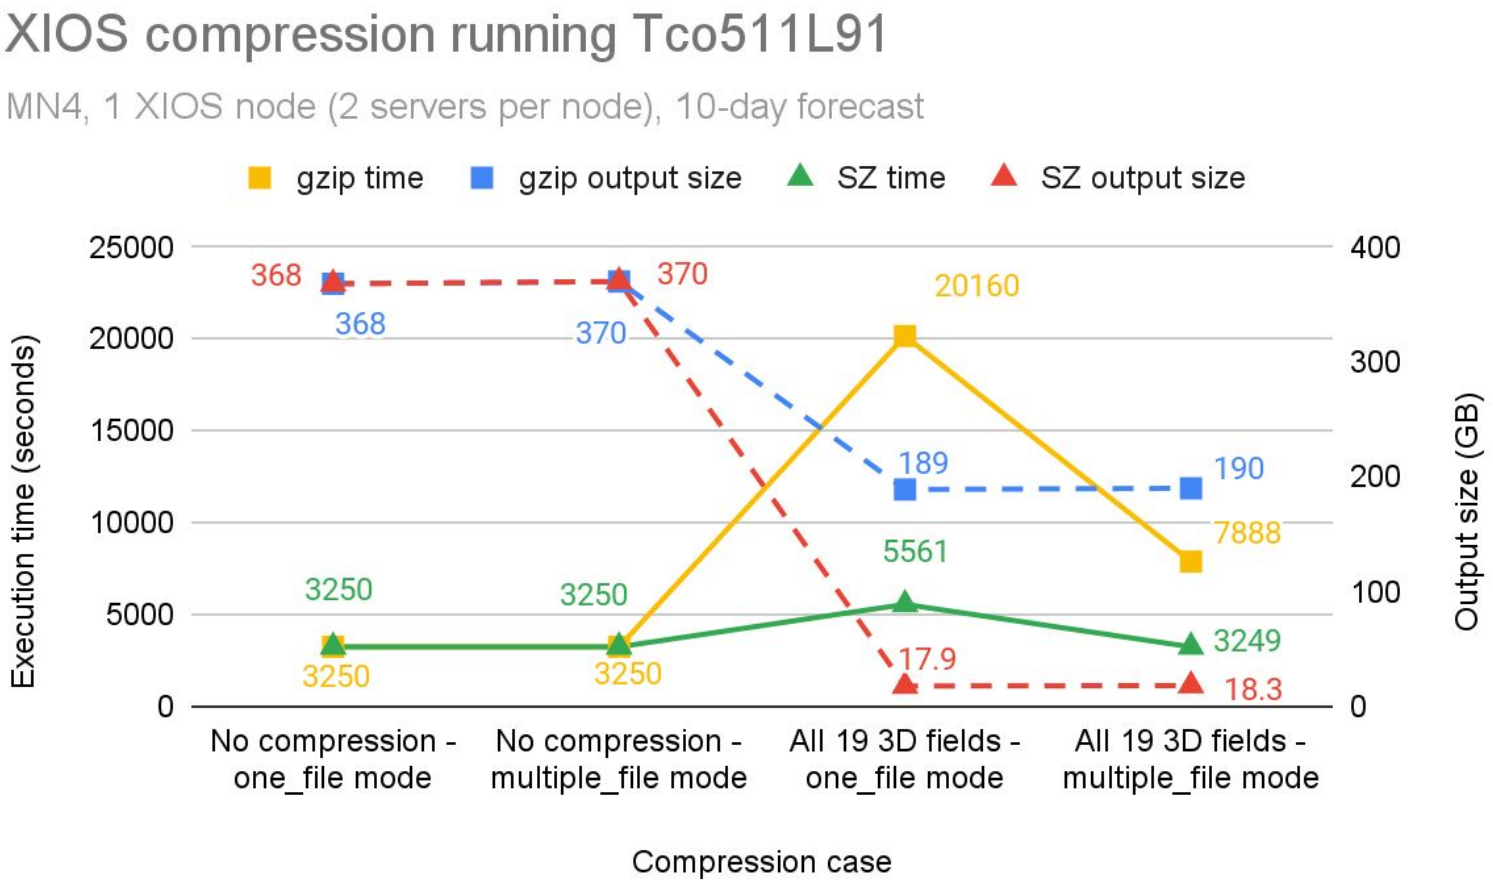

- gzip: compression level 5
- SZ: relative error bound 0.001

Conclusions and future work

- SZ lossy compressor is **faster** than the default gzip lossless compressor, achieving much **higher** compression ratios.
- We have now **complete control** of the **accuracy** of each field.
  - Determine the adequate compression parameters (error bounds) for each field depending on the acceptable errors.
- It is possible to use **compression** (both gzip and SZ filter) with **parallel I/O** in HDF5.
  - Test the latest version of HDF5 1.10.X or 1.12.X for performance improvements.
- These two tasks will be further explored and extended in these new European projects: DestinE CATS, EERIE and ESiWACE-3.

Open questions and collaboration opportunities

- Continuing the collaboration between ANL and BSC.
- Is it possible to specify the error bound in number of significant bits instead of decimal precision?
- Is it possible to use HDF5 1.12.X with the H5Z-SZ filter?
- Any feedback will be welcome.

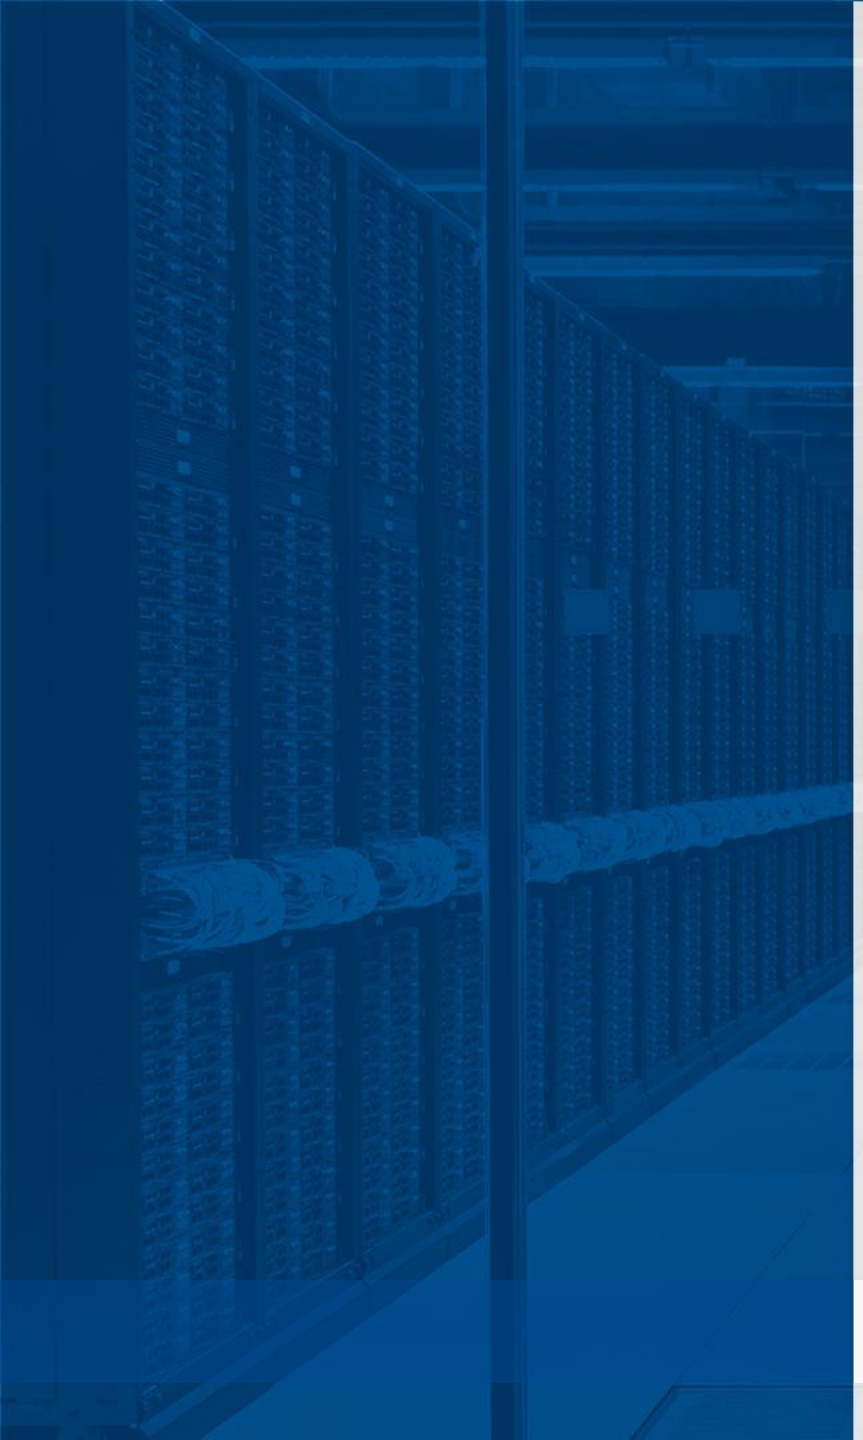

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thank you

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

xavier.yepes@bsc.es