Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Deploying the MONARCH atmospheric model in the Cloud

Miguel Castrillo, Francesca Macchia, Gilbert Montané, Kim Serradell

Computational Earth Sciences

09/02/2021

Weather and climate in the cloud

Outline

- Mineral dust services at BSC
- First experience (AWS)
- PoC with ORACLE
- Main takeaways
- Conclusions

Earth Sciences

Environmental modelling and forecasting, with a particular focus on weather, climate and air quality

Research and Forecasts

Service Users Sectors


Infrastructures


Solar Energy

Urban development


Transport


Wind Energy


Agriculture


Insurance

Mineral dust services at BSC

BSC dust operational forecast:

- Contribution to the SDS-WAS (regional) and ICAP (global) multi-model ensembles

WMO Dust Regional Centres:

- **Barcelona Dust Forecast Centre**. First specialized WMO Centre for mineral dust prediction. Started in 2014 - **Operational**
  - http://dust.aemet.es
  - @Dust_Barcelona
- Sand and Dust Storms Warning Advisory and Assessment System (**SDS-WAS**). North Africa, Middle East and Europe Regional Center. Started in 2010 – **Research**
  - http://sds-was.aemet.es
- Both WMO Regional Centres are jointly managed by BSC and AEMET

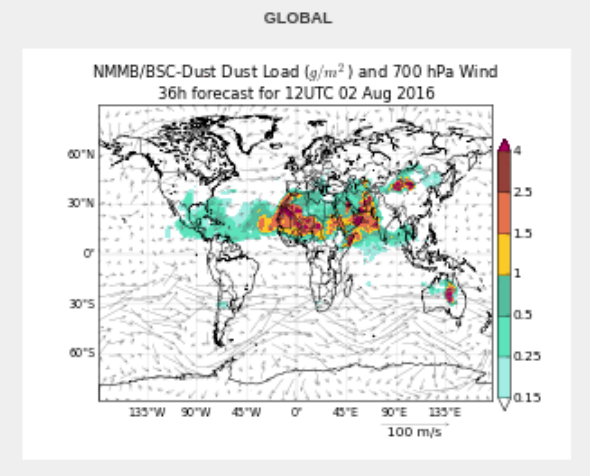

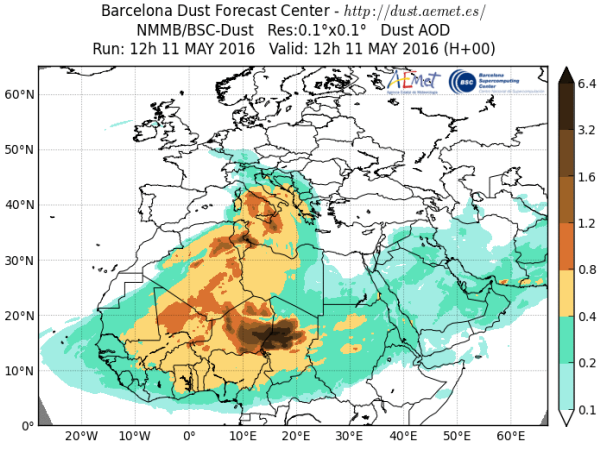

Operational services with some constraints

- BSC is a research center not an operational one. That means:
  - No 24/7 helpdesk for HPC services
  - No 24/7 forecast managers
  - [Annual] power maintenances
  - No HPC redundancy

- But…
  - we need to deliver our products.
  - Customers are always expecting the daily output (even if they know our constraints)

- Two solutions
  - Duplication of processes
    - BDFC simulation is executed in Mare Nostrum4 and Nimbus (AEMET)
  - Cloud

Why moving to the cloud?

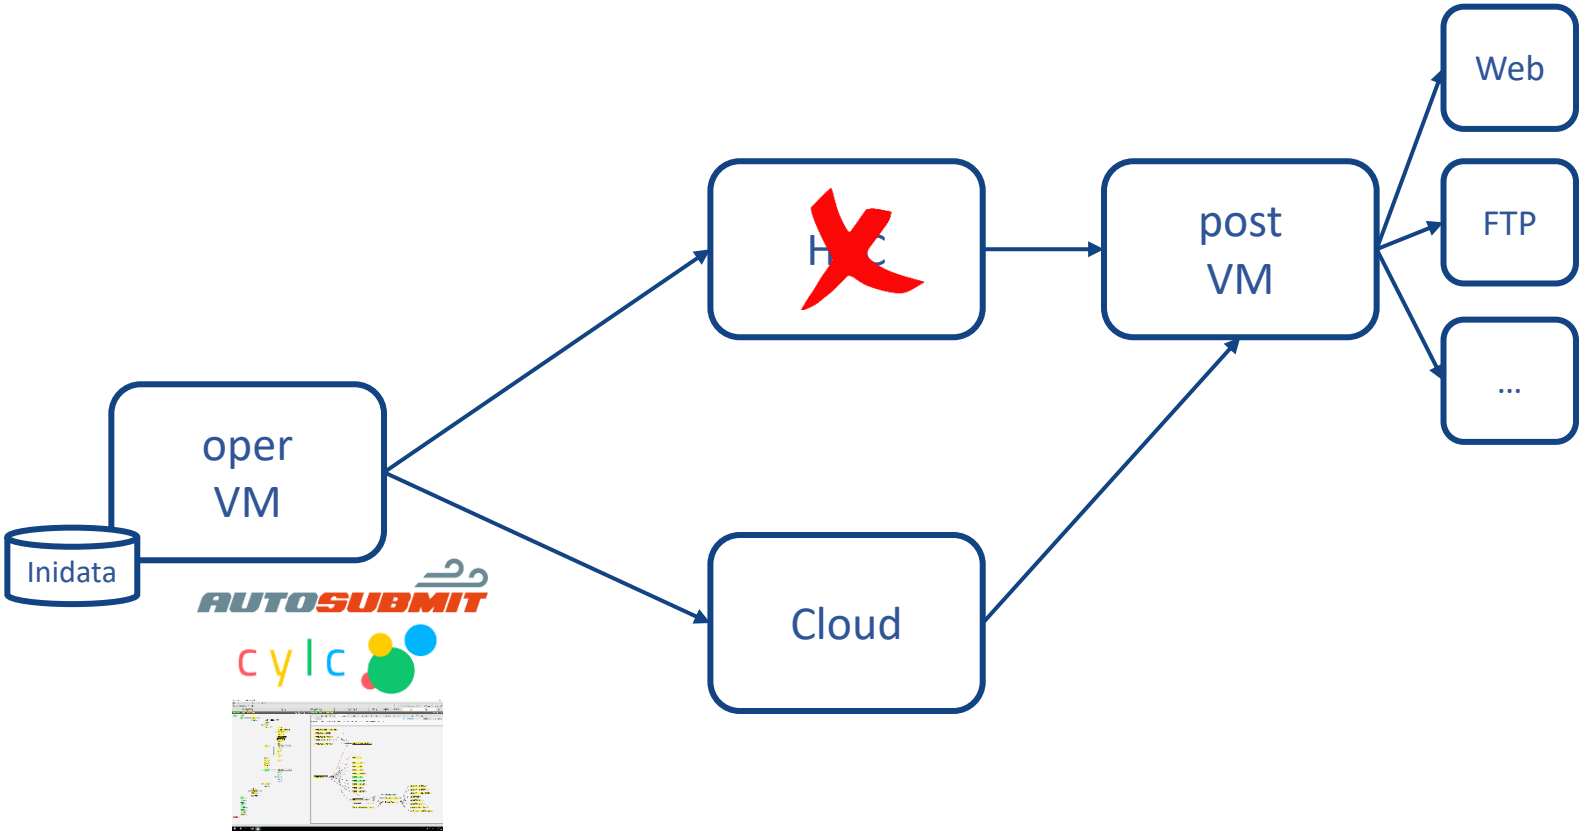

Amazon Web Services

- In 2019, we contacted AWS for a test
- We got some free credits
- Work done as final graduate work
  - Evaluation of AWS for running BSC-EC models, Adrià Quesada, 2019
- We deployed NMMB/MONARCH model (previous version)

- We did a **MANUAL** deployment
  - Starting from scratch from a AWS Ubuntu ISO
  - Install packages and/or compile dependencies
  - Build a small SLURM cluster
  - Compile model (Fortran+MPI)
  - Run model
  - Analyse results

Amazon Web Services (II)

- Results
  - The experience was successful
  - We validated the idea of running a forecast in the cloud
  - But everything was manual

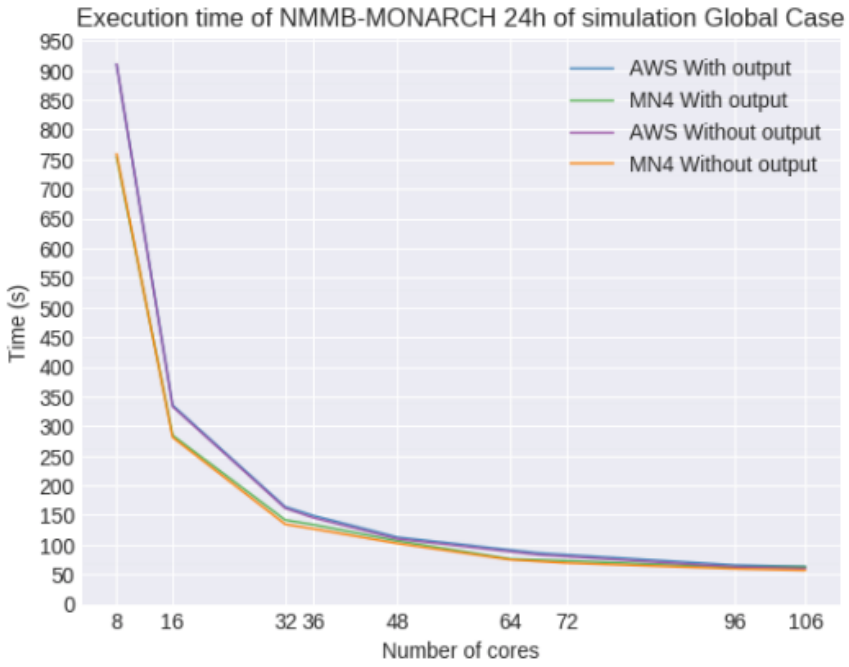

Amazon Web Services (III)

- Bonus track: extract traces using EXTRAE + PARAVER

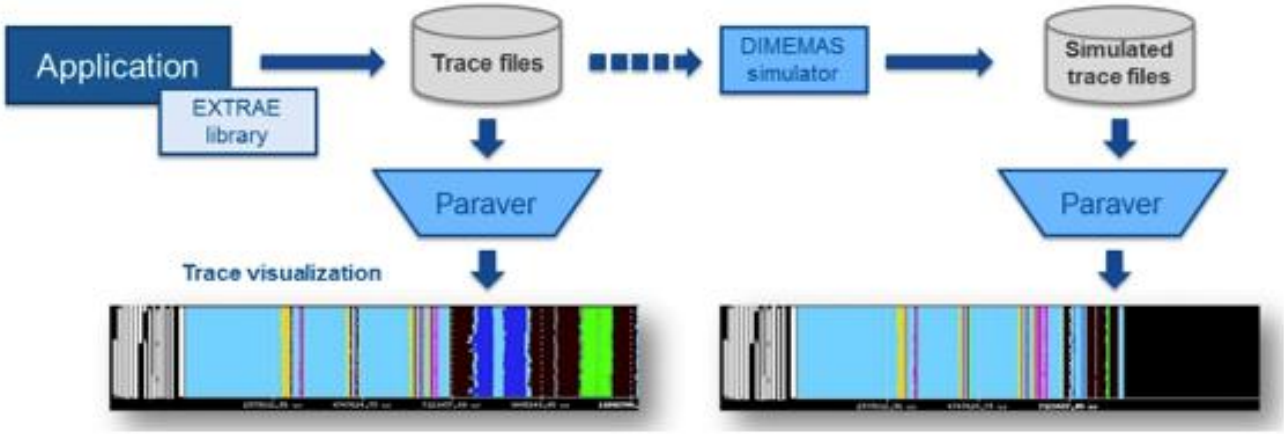

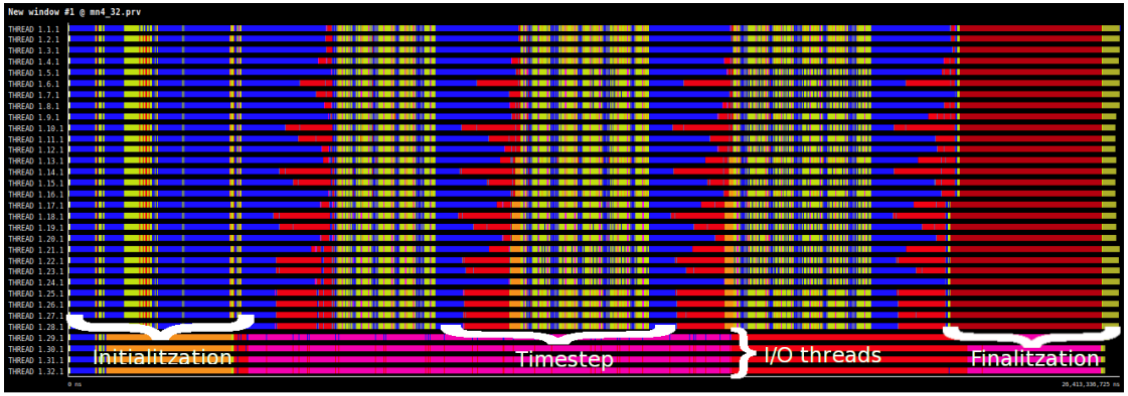

Figure 11: MN4 32 Cores Trace default view

Parallel tasks

- In the meantime, many efforts in parallel
  - Refactor of the code (from NMMB/MONARCH to MONARCH)
  - Development of Auto-MONARCH (code supported by Autosubmit, our workflow manager)
  - Gitlab + Jenkins integration
  - Output checker
  - Hiring of two Research Engineers

AUTOSUBMIT

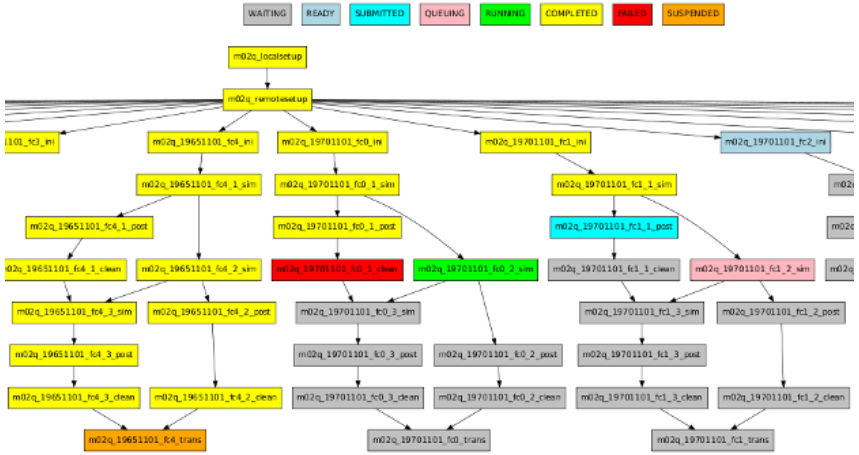

ORACLE

- BSC-ES has a partnership with HPCNow!
  - Spanish consultancy providing advanced supercomputing services for science and engineering.
  - They become partners with ORACLE Cloud.
  - Agreement between the three parts for a PoC.

- Three roles
  - Cloud provider (ORACLE Cloud)
  - HPC consultants (HPCNow!)
  - Research Engineers (Earth Sciences)

Workflow

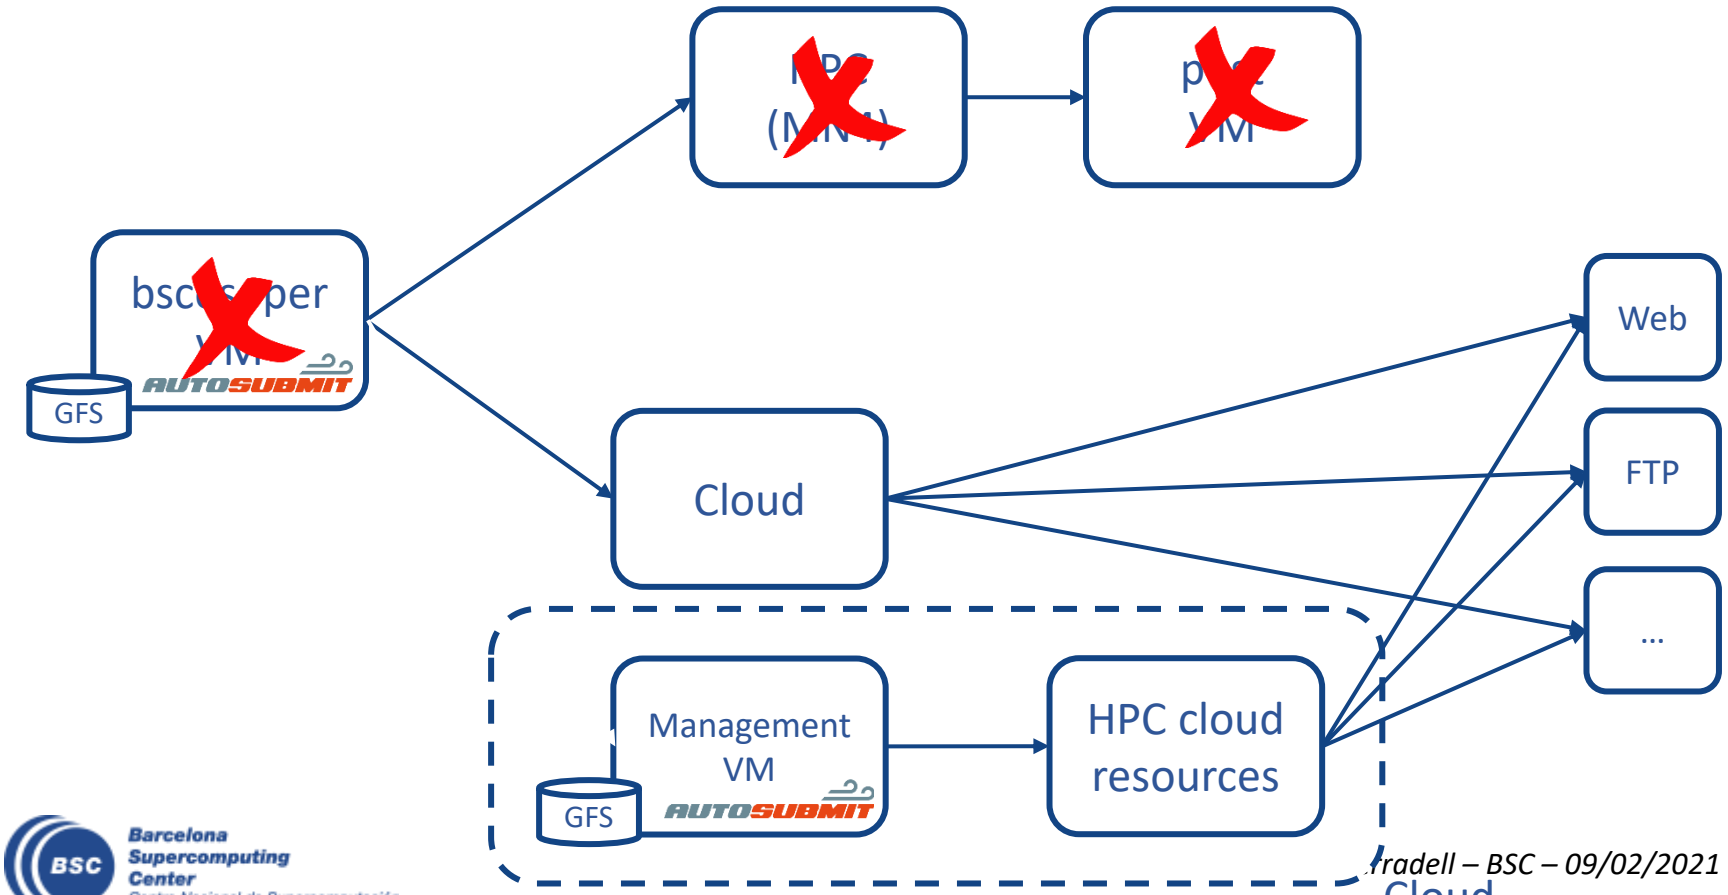
GFS
AUTOSUBMIT
Cloud
Web
FTP
...
Management VM
GFS
AUTOSUBMIT
HPC cloud resources
Cloud

Configuring HPC resources

- Between HPCNow! and ORACLE
  - Provision HPC resources. Configure the resources needed that will be available on request
    - Deploy a management node (login)
    - Using SLURM Cloud Scheduling capabilities
    - Sending a batch job will deploy automatically the resources
    - Efficient control of these resources

	Resource for the PoC
Management Node	1 x VM.Standard.E2.4
Storage	1 x VM.Standard.E2.4 + 10 VPU (Total: 6TB)
Compute 1 - AMD	2 x BM.Standard.E2.64
Compute 2 - Intel	2 x BM.HPC2.36

Software deployment

- Between BSC-ES and HPCNow!
  - Deploy the software stack
    - Using EasyBuild to automate the installation
    - Reproducibility and performance
    - Replicate your usual software stack


```
------------------------------------------------------------------------------------------ /storage/home/bsc02/.local/easybuild/modules/all ------------------------------------------------------------------------------------------
   Autoconf/2.69-GCCcore-9.3.0                Java/11.0.8-aarch64                (11)    ScaLAPACK/2.1.0-gompi-2020a            gettext/0.20.1                 (D)    libtool/2.4.6-GCCcore-9.3.0
   Automake/1.16.1-GCCcore-9.3.0              LLVM/9.0.1-GCCcore-9.3.0                   Szip/2.1.1-GCCcore-9.3.0               gompi/2020a                           libunwind/1.3.1-GCCcore-9.3.0
   Autotools/20180311-GCCcore-9.3.0           LibTIFF/4.1.0-GCCcore-9.3.0                Tcl/8.6.10-GCCcore-9.3.0               gperf/3.1-GCCcore-9.3.0               libxml2/2.9.10-GCCcore-9.3.0
   Bison/3.3.2                                LittleCMS/2.9-GCCcore-9.3.0                Tk/8.6.10-GCCcore-9.3.0                groff/1.22.4-GCCcore-9.3.0            lz4/1.9.2-GCCcore-9.3.0
   Bison/3.5.3-GCCcore-9.3.0                  M4/1.4.18-GCCcore-9.3.0                    UCX/1.8.0-GCCcore-9.3.0                gzip/1.10-GCCcore-9.3.0               makeinfo/6.7-GCCcore-9.3.0
   Bison/3.5.3                         (D)    M4/1.4.18                          (D)     UnZip/6.0-GCCcore-9.3.0                help2man/1.47.4                       ncurses/6.1
   CMake/3.16.4-GCCcore-9.3.0                 Mako/1.1.2-GCCcore-9.3.0                   X11/20200222-GCCcore-9.3.0             help2man/1.47.12-GCCcore-9.3.0 (D)    ncurses/6.2-GCCcore-9.3.0      (D)
   DB/18.1.32-GCCcore-9.3.0                   Mesa/20.0.2-GCCcore-9.3.0                  XZ/5.2.5-GCCcore-9.3.0                 hwloc/2.2.0-GCCcore-9.3.0             netCDF-Fortran/4.5.2-gompi-2020a
   Doxygen/1.8.17-GCCcore-9.3.0               Meson/0.55.1-GCCcore-9.3.0-Python-3.8.2    Xvfb/1.20.9-GCCcore-9.3.0              intltool/0.51.0-GCCcore-9.3.0         netCDF/4.7.4-gompi-2020a
   FFTW/3.3.8-gompi-2020a                     NASM/2.14.02-GCCcore-9.3.0                 binutils/2.34-GCCcore-9.3.0            libGLU/9.0.1-GCCcore-9.3.0            nettle/3.6-GCCcore-9.3.0
   GCC/9.3.0                                  NLopt/2.6.1-GCCcore-9.3.0                  binutils/2.34                   (D)    libdrm/2.4.100-GCCcore-9.3.0          numactl/2.0.13-GCCcore-9.3.0
   GCCcore/9.3.0                              Ninja/1.10.0-GCCcore-9.3.0                 bzip2/1.0.8-GCCcore-9.3.0              libevent/2.1.11-GCCcore-9.3.0         pixman/0.38.4-GCCcore-9.3.0
   GLPK/4.65-GCCcore-9.3.0                    OpenBLAS/0.3.9-GCC-9.3.0                   cURL/7.69.1-GCCcore-9.3.0              libfabric/1.11.0-GCCcore-9.3.0        pkg-config/0.29.2-GCCcore-9.3.0
   GLib/2.64.1-GCCcore-9.3.0                  OpenMPI/4.0.3-GCC-9.3.0                    cairo/1.16.0-GCCcore-9.3.0             libffi/3.3-GCCcore-9.3.0              util-linux/2.35-GCCcore-9.3.0
   GMP/6.2.0-GCCcore-9.3.0                    PCRE/8.44-GCCcore-9.3.0                    expat/2.2.9-GCCcore-9.3.0              libglvnd/1.2.0-GCCcore-9.3.0          xorg-macros/1.19.2-GCCcore-9.3.0
   GSL/2.6-GCC-9.3.0                          PCRE2/10.34-GCCcore-9.3.0                  flex/2.6.4-GCCcore-9.3.0               libiconv/1.16-GCCcore-9.3.0           zlib/1.2.11-GCCcore-9.3.0
   Ghostscript/9.52-GCCcore-9.3.0             PMIx/3.1.5-GCCcore-9.3.0                   flex/2.6.4                      (D)    libjpeg-turbo/2.0.4-GCCcore-9.3.0     zlib/1.2.11                    (D)
   HDF5/1.10.6-gompi-2020a                    Perl/5.30.2-GCCcore-9.3.0                  fontconfig/2.13.92-GCCcore-9.3.0       libpciaccess/0.16-GCCcore-9.3.0       zstd/1.4.4-GCCcore-9.3.0
   ICU/66.1-GCCcore-9.3.0                     Python/2.7.18-GCCcore-9.3.0                foss/2020a                             libpng/1.6.37-GCCcore-9.3.0
   ImageMagick/7.0.10-1-GCCcore-9.3.0         Python/3.8.2-GCCcore-9.3.0         (D)     freetype/2.10.1-GCCcore-9.3.0          libreadline/8.0-GCCcore-9.3.0
   JasPer/2.0.14-GCCcore-9.3.0                SQLite/3.31.1-GCCcore-9.3.0                gettext/0.20.1-GCCcore-9.3.0           libsndfile/1.0.28-GCCcore-9.3.0
```

  - Compile and optimize your model to get the best performance
    - You don't know the hardware

Auto-MONARCH - Configuration

- MONARCH NAME-E Regional
  - Res: 0.1º x 0.1º
  - Grid: 1021 x 701 x 40
  - 24h Forecast
  - Scaling exercise: 68, 144, 288 and 576 cores.

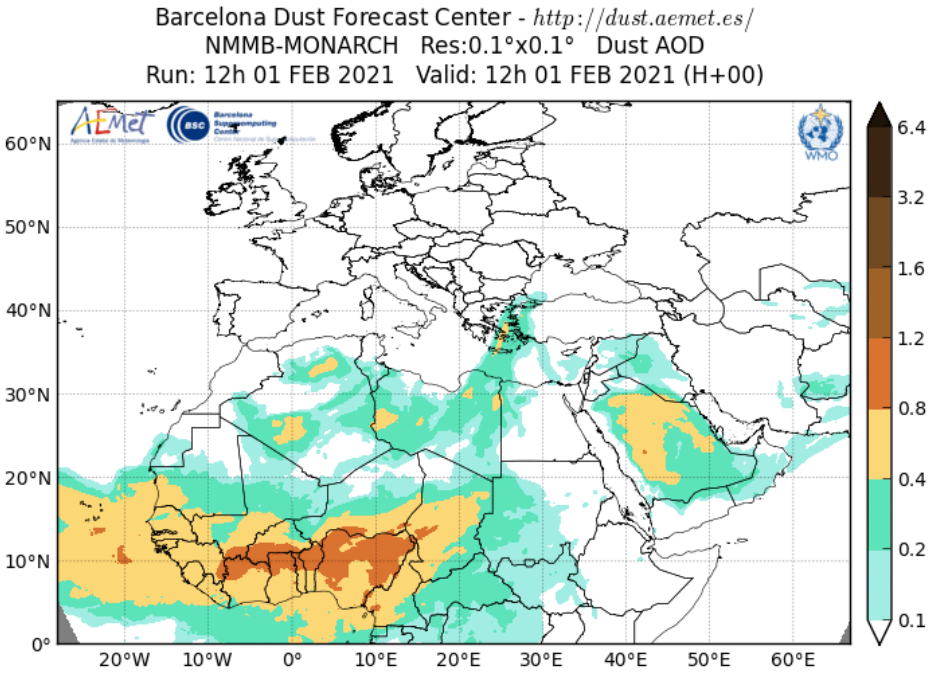

Auto-MONARCH - Results

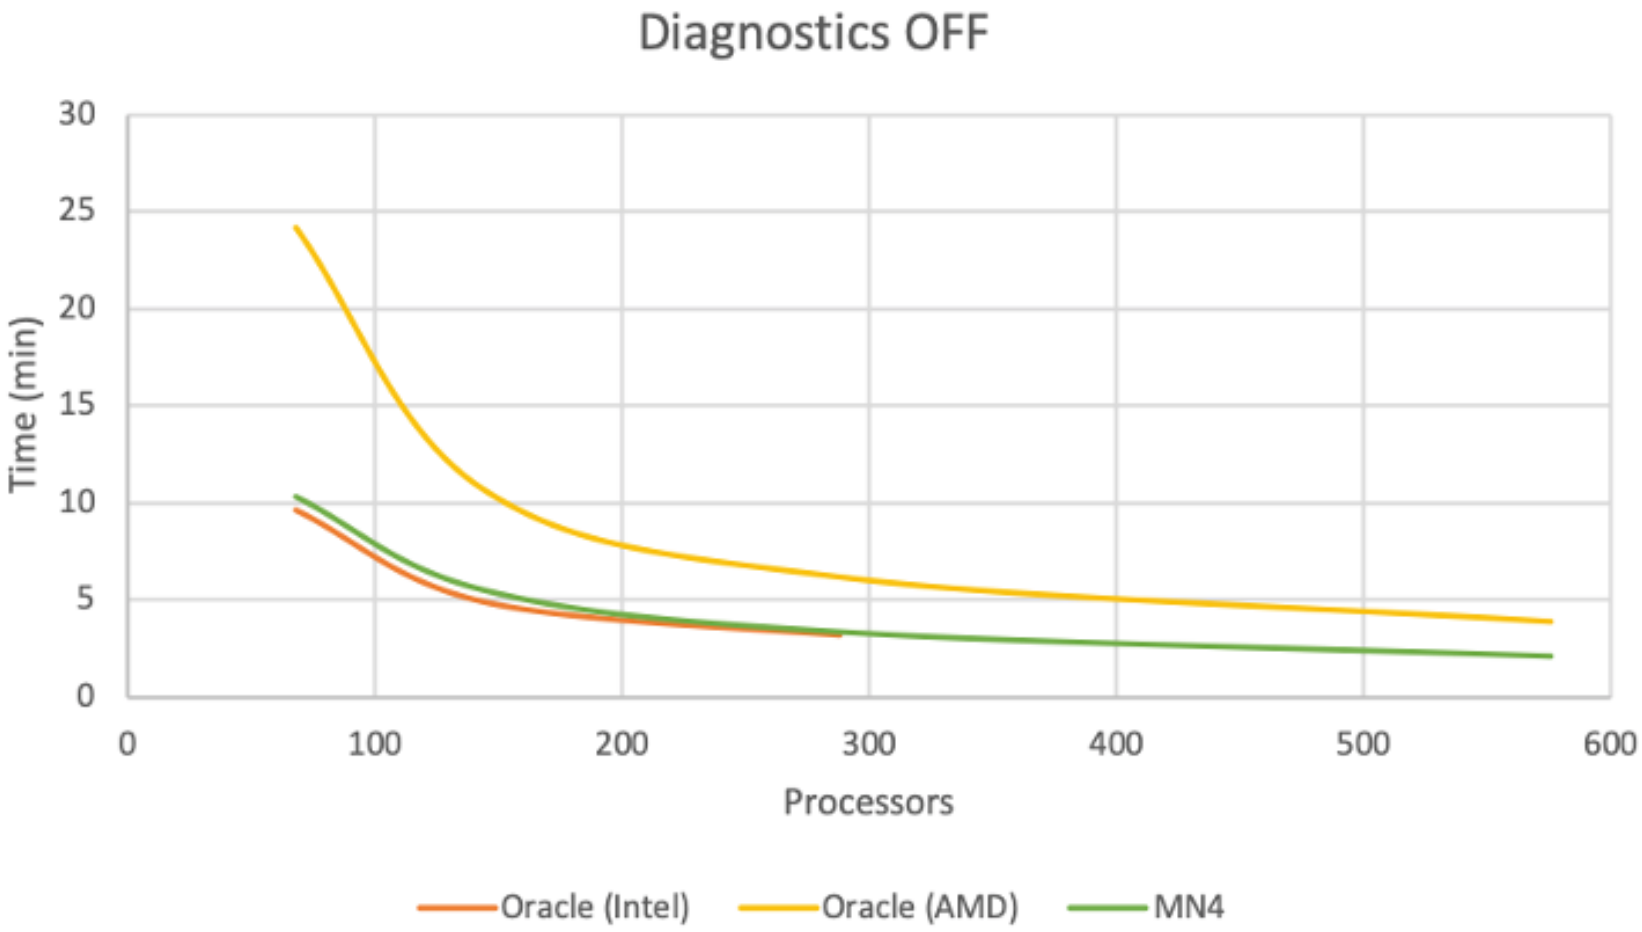

Main Takeaways

- The cloud is a **feasible solution** to replace usual HPC premises as a backup solution
- **Collaboration** with vendors is mandatory
- **Automate** your processes as much as possible
- **Replicate** the software stack is key
- **Verify** your results
- Invest on **human resources**
- And Price?

Future work

- Deploy operationally this solution
- Develop “smart” capabilities for Autosubmit
- Extend this approach to other BSC forecasts

Acknowledgements

- ORACLE
- HPCNow!
- Computational Earth Sciences research engineers and scientists from Atmospheric Composition group and Earth System Services group.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Thank you

kim.serradell@bsc.es