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

Continuation of ESCAPE-2 efforts in ESiWACE2

Florian Ziemen (DKRZ) on behalf of Joachim Biercamp
(who's on a cycling trip somewhere near the alps)

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

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

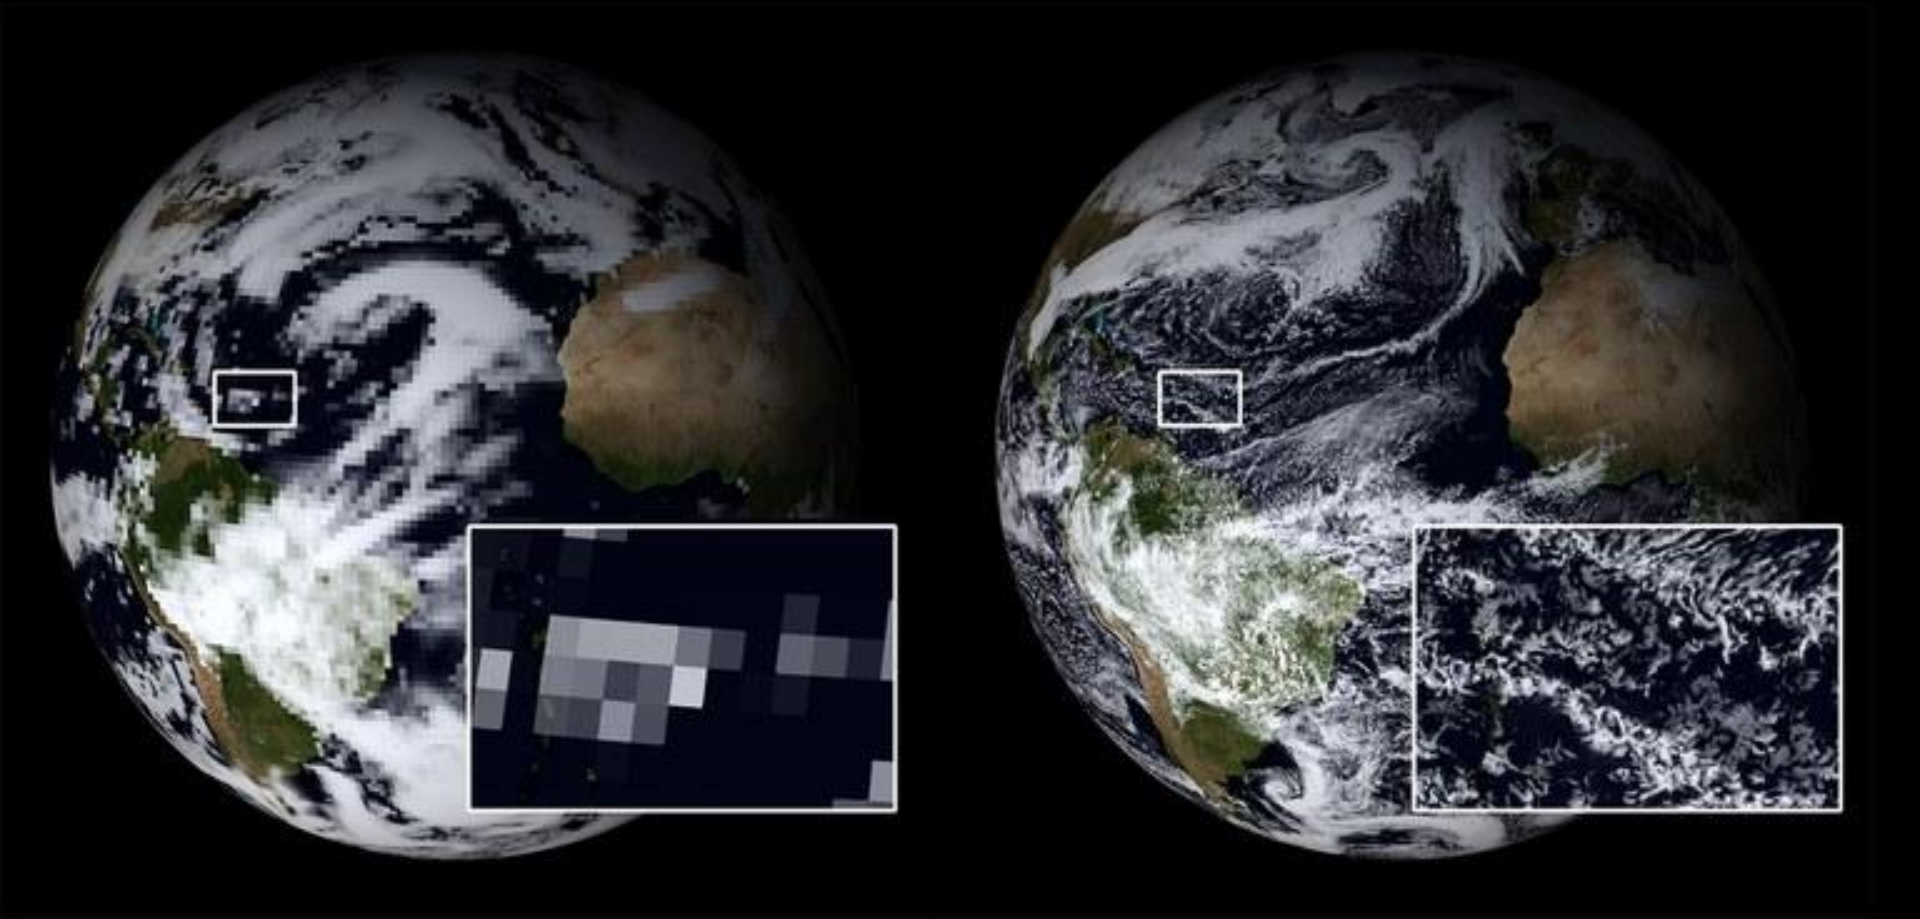

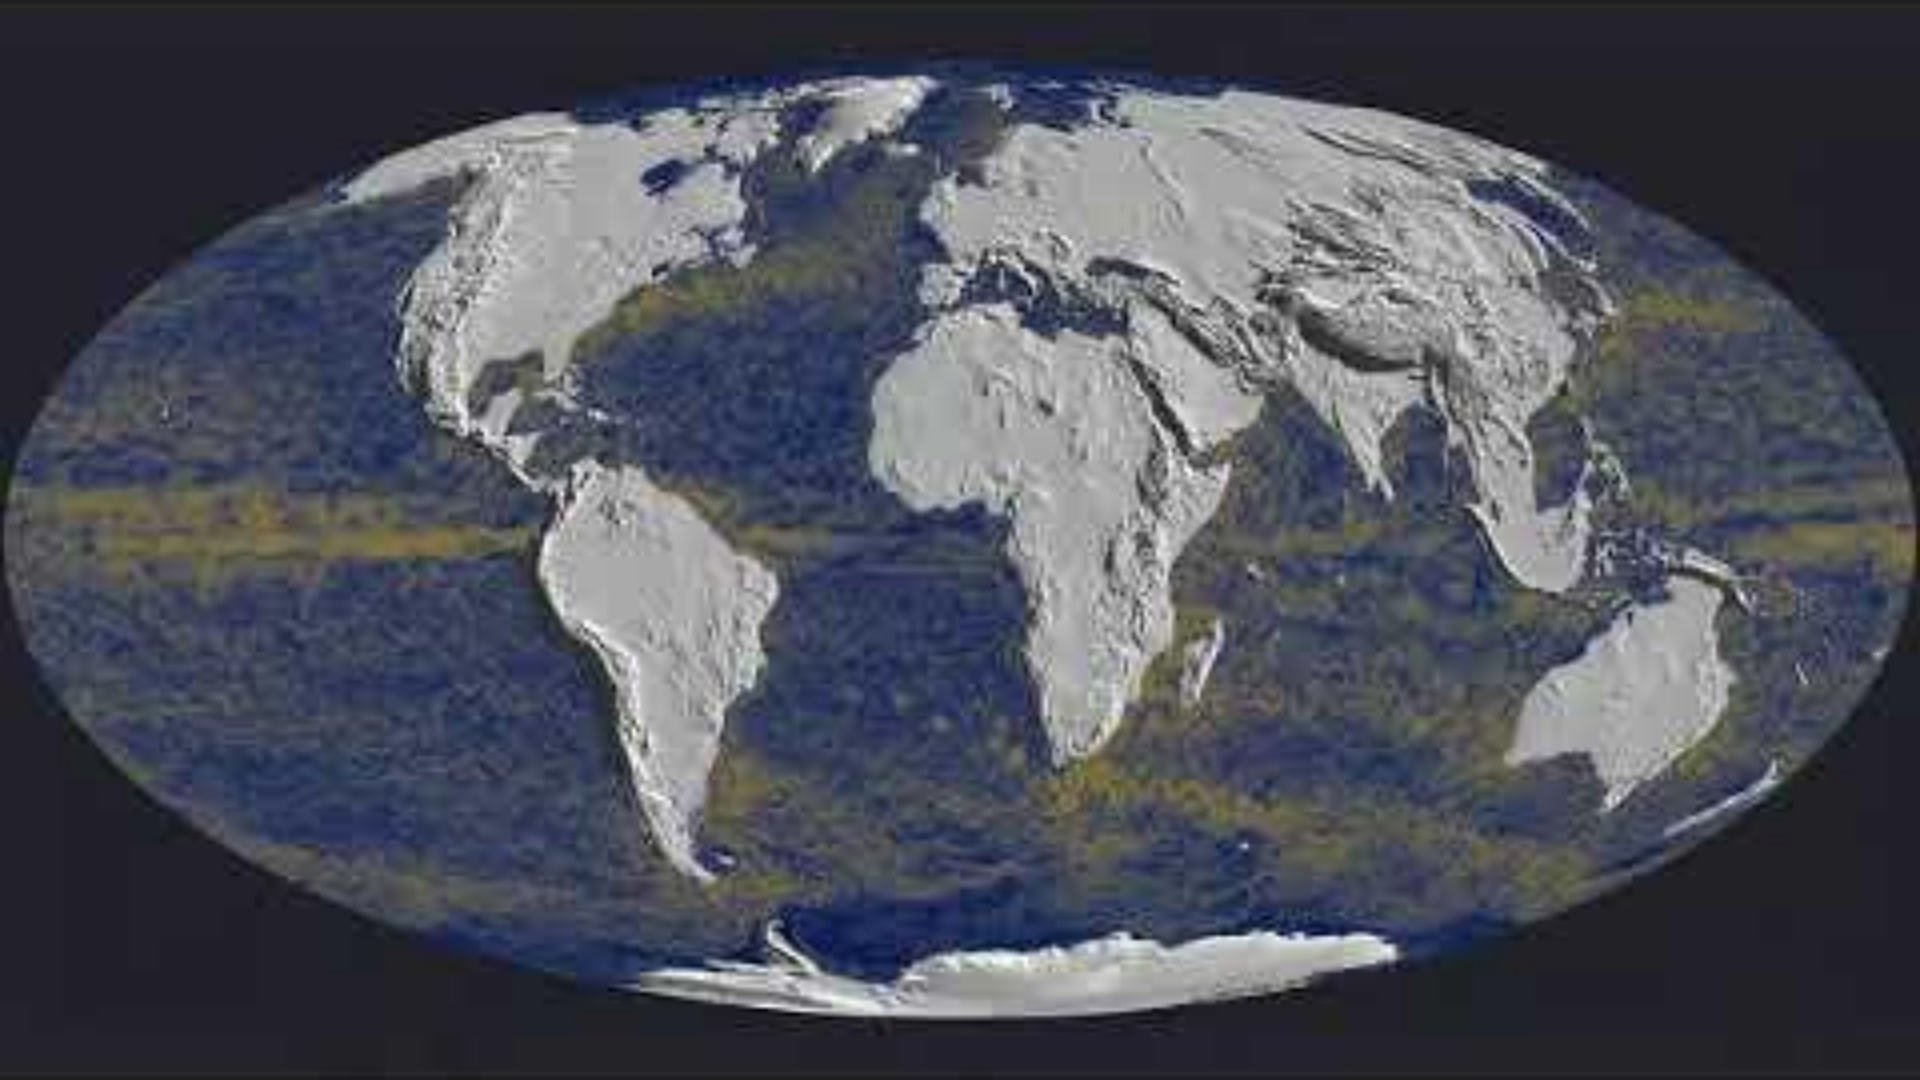

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ESiWACE2 is helping climate science

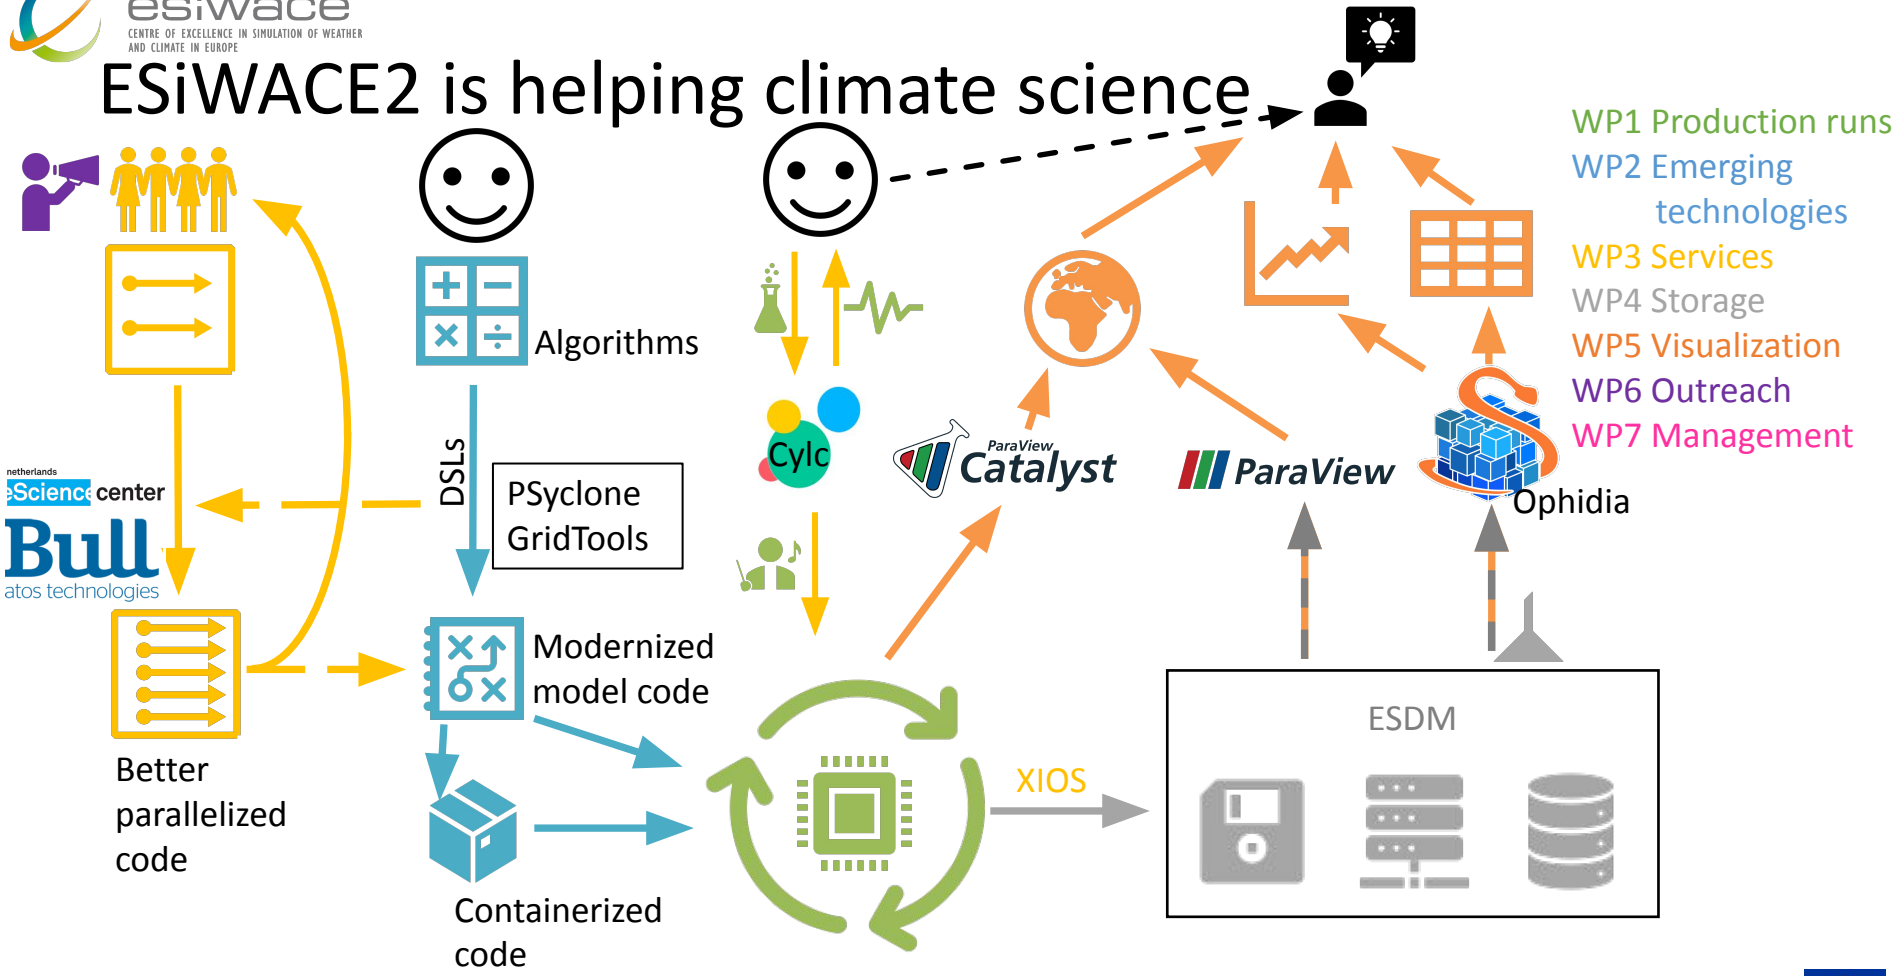

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

Cross-disciplinarity:

Partners from Weather, Climate, HPC, Industry

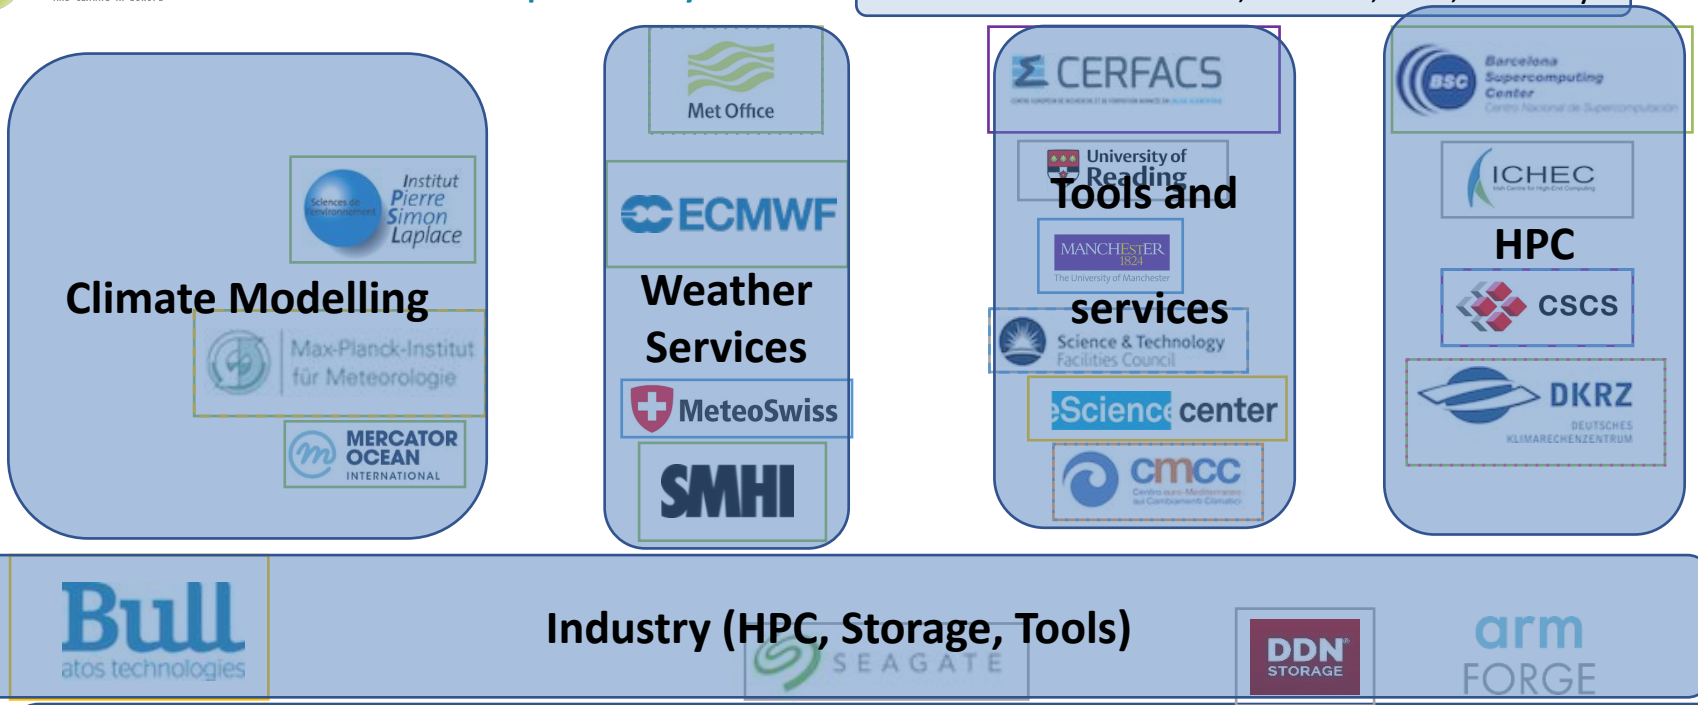

BULL optimization of NEMO for improving performance scalability at high resolution.

Seagate development of the Clovis/Mero Object store backend for ESDM.

DDN development of the infinity memory engine backend for ESDM.

WP1 Production runs WP2 Emerging technologies WP3 Services WP4 Storage WP5 Visualization
WP6 Outreach WP7 Management

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

HPCW Benchmarks

Partners involved: DKRZ (lead), BSC, BULL, ECMWF, ETH Zurich / CSCS

Extend the benchmarks

- GPUs if supported by models
- Additional experiments (needs input from the models' teams)
- Enhance Kronos for simulating the production mode of storm resolving models on exascale systems

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

HPCW Benchmarks

Tackle hardware and compilation uncertainties by running on different Platforms

- DRKZ: Mistral, Levante (once installed)
- ECMWF
- upcoming EuroHPC pre-exascale systems
  - Lumi
  - MareNostrum 5
- …

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

DSLs

ESCAPE

ESiWACE

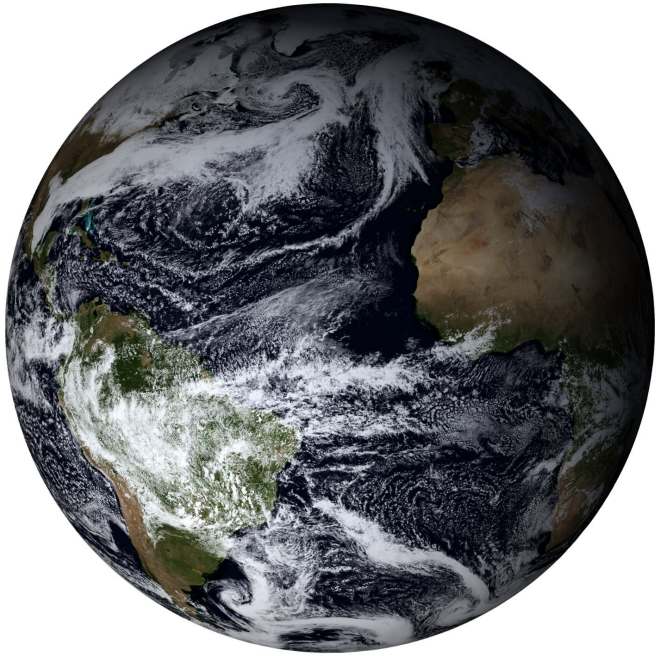

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

From ESCAPE-2 dwarfs to full models

- ESCAPE-2 developed a complete DSL toolchain for weather and climate models. Demonstrated in ESCAPE2 dwarfs: NEMO, ICON-O.

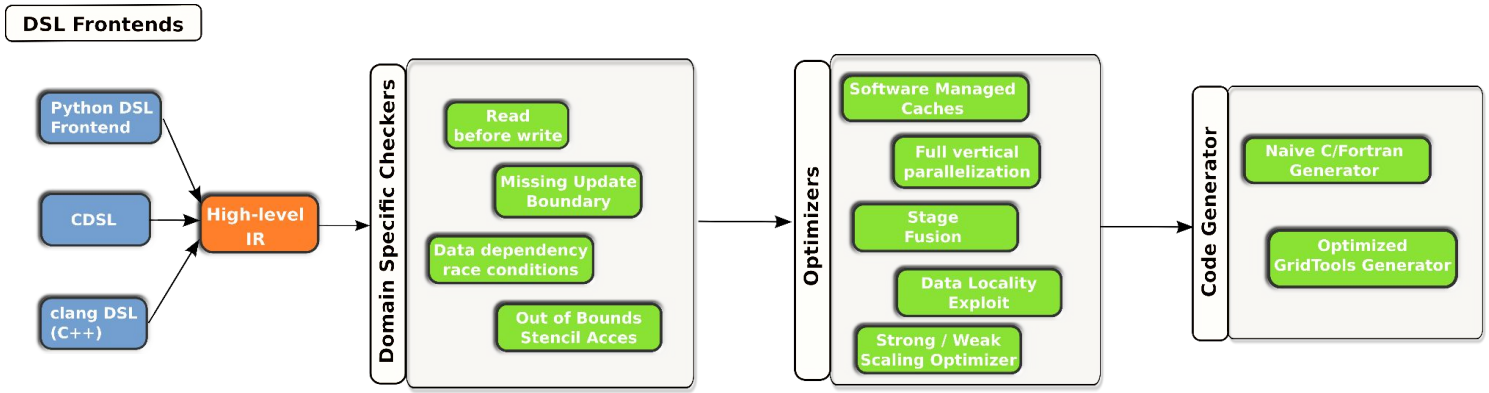

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

From ESCAPE-2 dwarfs to full models

- ESCAPE-2 developed a complete DSL toolchain for weather and climate models. Demonstrated in ESCAPE2 dwarfs: NEMO, ICON-O.
- ESiWACE-2 adopted the toolchain with the dusk python frontend, specifically designed for ICON (unstructured nature on icosahedral grids)
- An entire dry-dynamical core of ICON was implemented on the python DSL, integrated in operational model, NWP verifying on GPU

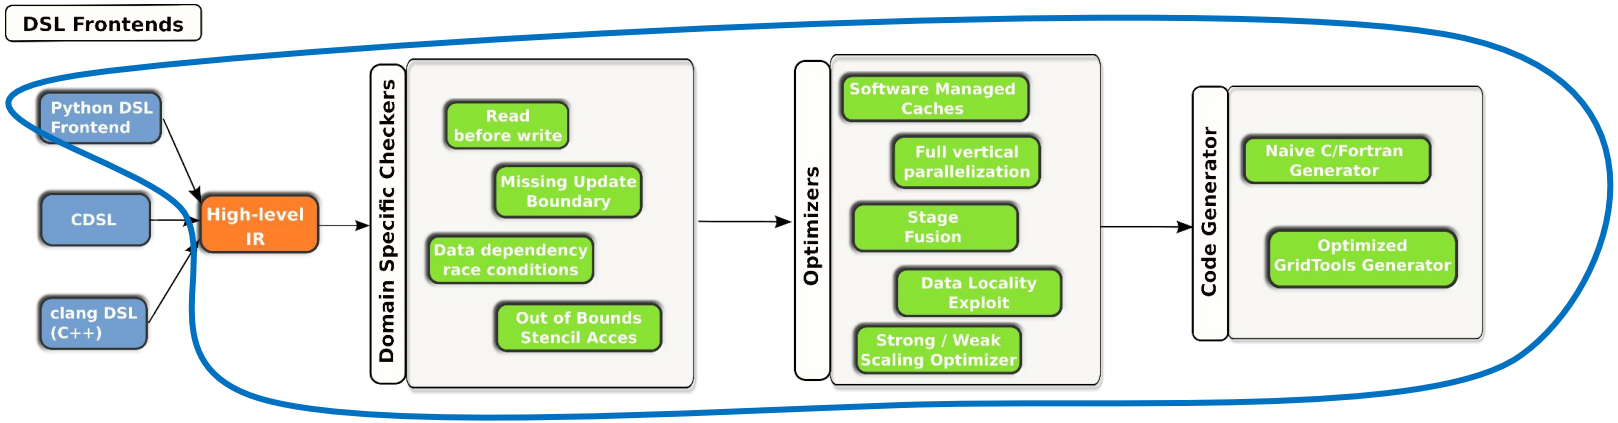

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

Example of ICON Fortran -> Python DSL translation

```
rl_start = start_bdydiff_e
rl_end   = grf_bdywidth_e

i_startblk = p_patch%edges%start_block(rl_start)
i_endblk   = p_patch%edges%end_block(rl_end)

...
```

```
        ! Lateral boundary diffusion for vn
     i_startblk = p_patch%edges%start_block(start_bdydiff_e)
     i_endblk   = p_patch%edges%end_block(grf_bdywidth_e)

!$OMP DO PRIVATE(je,jk,jb,i_startidx,i_endidx) ICON_OMP_DEFAULT_SCHEDULE
     DO jb = i_startblk,i_endblk

       CALL get_indices_e(p_patch, jb, i_startblk, i_endblk, &
                          i_startidx, i_endidx, start_bdydiff_e, grf_bdywidth_e)

!$ACC PARALLEL LOOP DEFAULT(NONE) GANG VECTOR COLLAPSE(2) ASYNC(1) IF( i_am_accel_node .AND. acc_on
)
       DO jk = 1, nlev
!DIR$ IVDEP
         DO je = i_startidx, i_endidx
           p_nh_prog%vn(je,jk,jb) =   &
             p_nh_prog%vn(je,jk,jb) + &
             z_nabla2_e(je,jk,jb) * &
             p_patch%edges%area_edge(je,jb)*fac_bdydiff_v
         ENDDO
       ENDDO
     ENDDO
!$OMP END DO
```

```
fac_bdydiff_v = Global("fac_bdydiff_v")

@stencil
def mo_nh_diffusion_stencil_09(
    z_nabla2_e: Field[Edge, K],
    area_edge: Field[Edge],
    p_nh_prog_vn: Field[Edge, K]
):
    with domain.upward.across[lb+1:nudging-1]:
        p_nh_prog_vn += z_nabla2_e * area_edge *
fac_bdydiff_v
```

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

Integration of the DSL into ICON Fortran

- # computation kernels

  velocity_adv: 22
  diffusion: 18
  solve_nonhydro: 60

  were translated to DSL and integrated in Fortran

- Still in Fortran
  - MPI Communication
  - Control flow

```
CALL sync_patch_array (SYNC_E, p_patch, p_nh_prog%vn,opt_varname="diffusion: vn sync"

...

IF (diffu_type == 3) THEN ! Only Smagorinsky diffusion
 IF ( jg == 1 .AND. l_limited_area .OR. jg > 1 .AND. .NOT. lfeedback(jg)) THEN
  ...
 ENDIF
ENDIF
```

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

Diffusion module translation - Performance

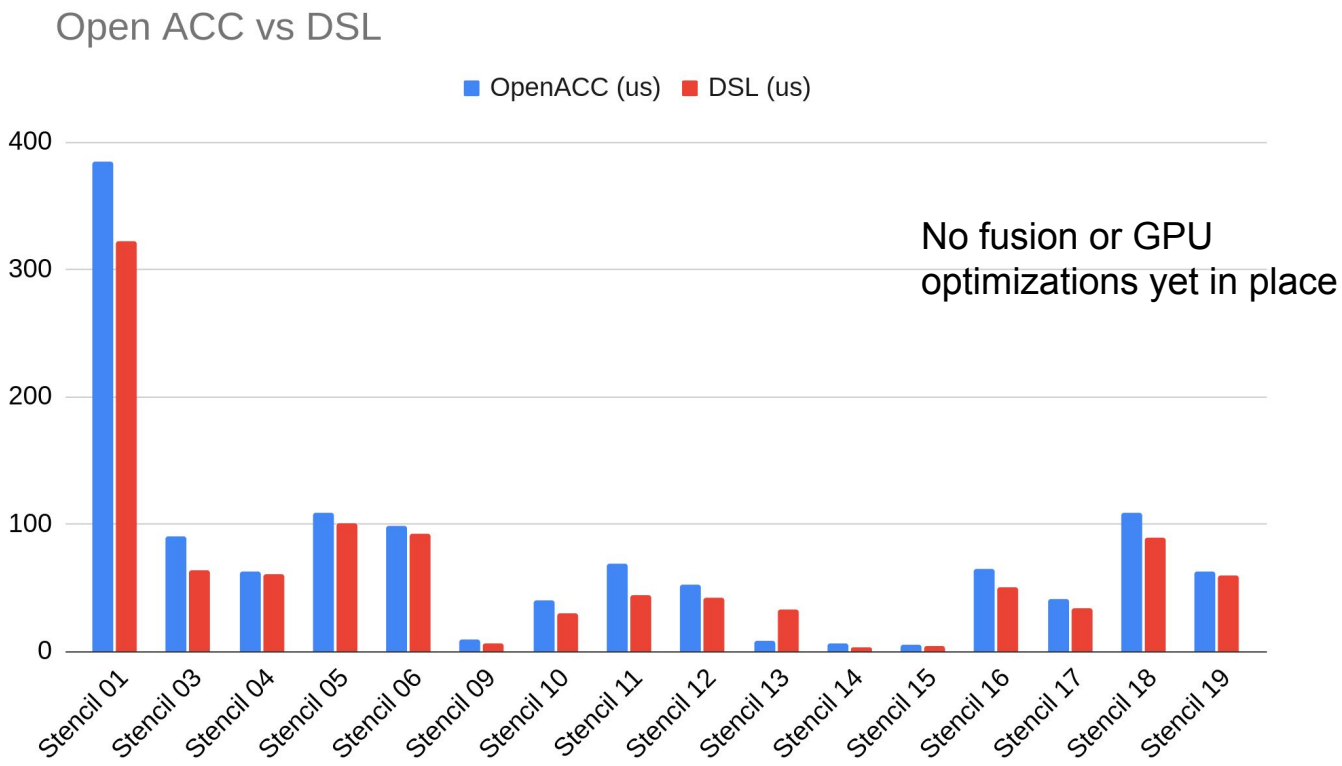

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

ESiWACE2 DSL Overview

- PSyclone and DAWN

- Demonstrators
  - ICON, IFS
  - NEMO, LFRic
- Comparison
  - NEMO Dwarf
  - PSyclone -> DAWN interoperability
- Training and dissemination
  - DSL training event
  - ESiWACE summer schools

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

LFRic Matvec Benchmark

- Full model runs with MPI and OpenMP
- Working on OpenACC version
- Matvec responsible for 30% of model computational cost
- Manual restructure matvec to improve performance
- OpenMP performance 2x faster than original
- OpenACC implementation performs well with a large number of columns

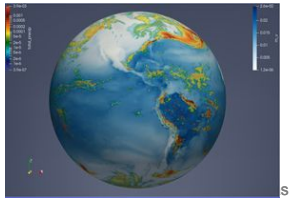

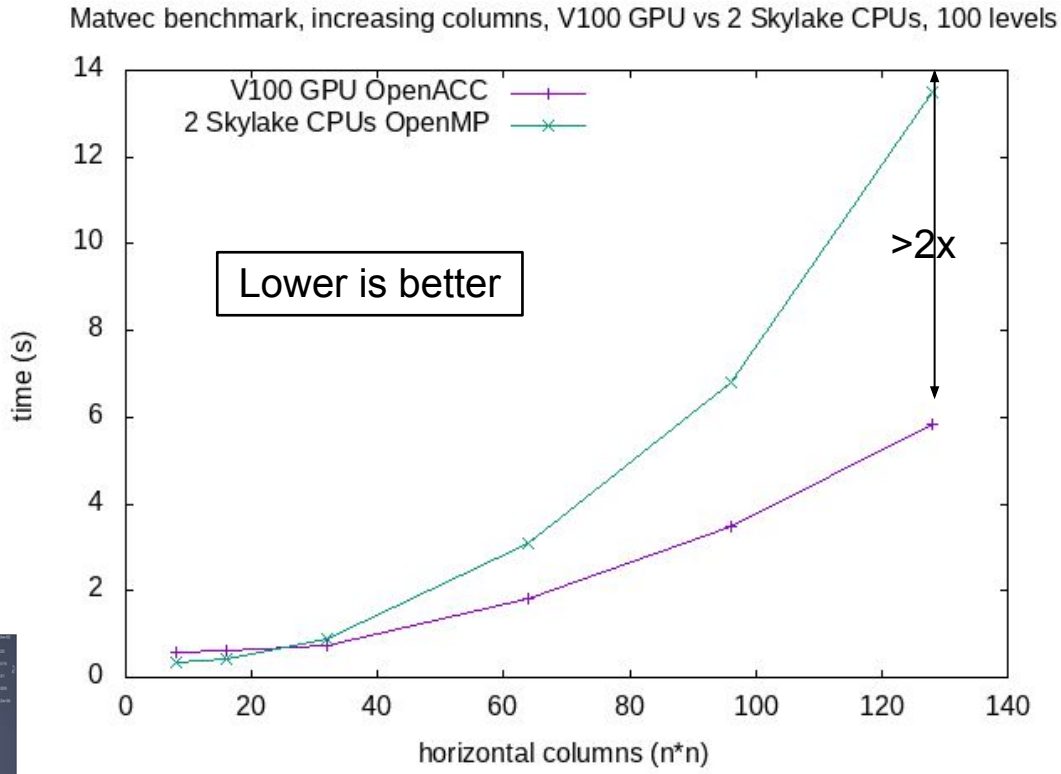

s received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

NEMO

- Full model runs on GPUs
- Concentrated on performance without sea-ice first
- ORCA1 Single GPU
- Multi-GPU ORCA12 running on Marconi but poor scaling. About to test on Juwels Booster.
- Start sea-ice optimisation next. Currently 75% time on GPU (cf 11% on CPU).

"Too good to be true" – Andreas Mueller ECMWF 2019

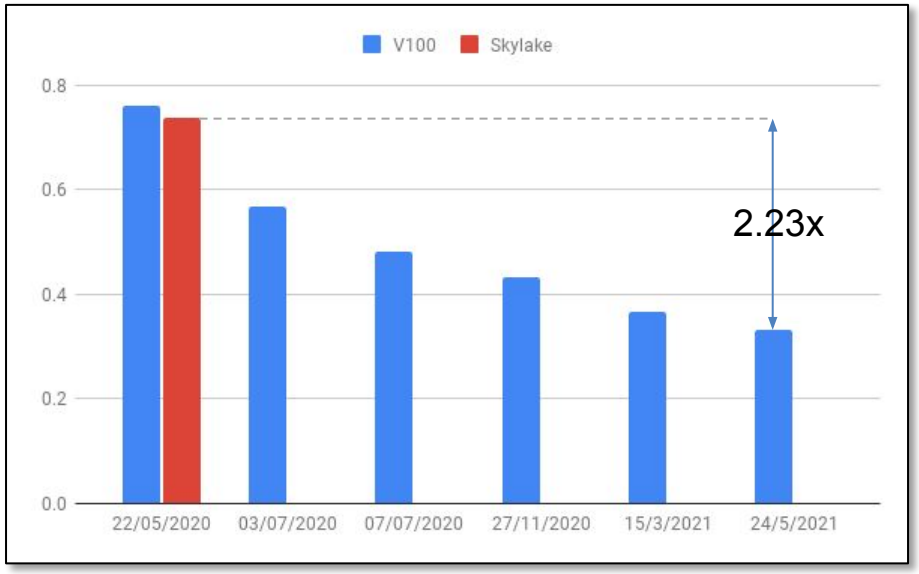

Lower is better

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

DSL Interoperability

- Different DSL front ends (and probably back ends)
- Compatible IRs?
- Evaluate PSyclone to DAWN via PSyIR and SIR

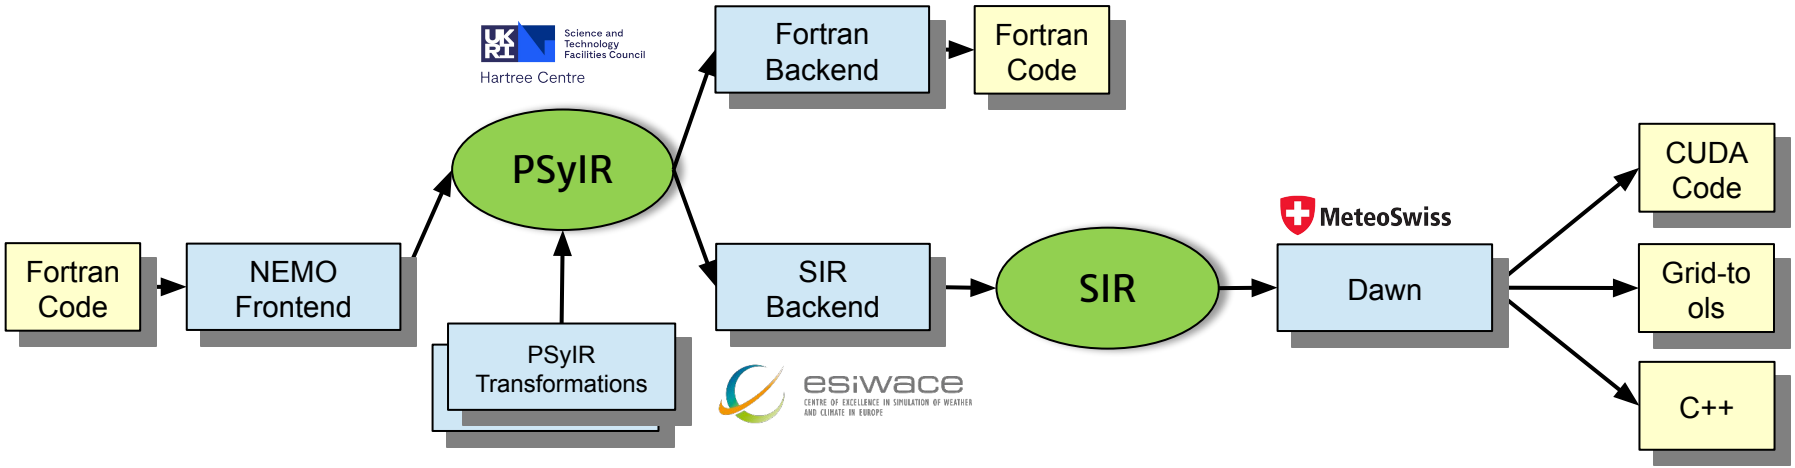

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

DSL Interoperability

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

- NEMO Dwarf – tracer advection
- PSyclone transforms
  - Intrinsics
    - ✔ABS, SIGN, MIN
  - Perfectly-nested triple loops
    - ✔Array assignment -> loops
    - ✔Single array index -> loops
    - ⚠ Loop invariant scalars
  - ☐ 2D arrays

```
program hori_diff
  do k=1,n
     do j=1,n
        do i=1,n
           lap(i,j,k)=(-4.0)*in(i,j,k)+coeff(i,j,k)*( &
                      in(i+1,j,k)+in(i-1,j,k)+in(i,j+1,k)+in(i,j-1,k))
           out(i,j,k)=(-4.0)*lap(i,j,k)+coeff(i,j,k)*( &
                      lap(i+1,j,k)+lap(i-1,j,k)+lap(i,j+1,k)+lap(i,j-1,k))
        end do
     end do
  end do
end program hori_diff
```

```
 // initialized iterators
 int idx111 = (blockIdx.x*32+iblock)*1+(blockIdx.y*4+jb      )*stride_111_1;

 // jump iterators to match the intersection of beginning of next interval and the parallel execution block
 idx111 += max(0, blockIdx.z * 4) * stride_111_2;
 int kleg_lower_bound = max(0,blockIdx.z*4);
 int kleg_upper_bound = min( ksize - 1 + 0,(blockIdx.z+1)*4-1);;
for(int k = kleg_lower_bound+0; k <= kleg_upper_bound+0; ++k) {
 if(iblock >= -1 && iblock <= block_size_i -1 + 1 && jblock >= -1 && jblock <= block_size_j -1 + 1) {
lap[idx111] = (((gridtools::clang::float_type) -4.0 * __ldg(&(in[idx111]))) + (__ldg(&(coeff[idx111])) * (((__ldg(&
(in[idx111+1*1])) + __ldg(&(in[idx111+1*-1]))) + __ldg(&(in[idx111+stride_111_1*1]))) + __ldg(&(in[idx111+stride_11
1_1*-1])))));
 }    __syncthreads();
 if(iblock >= 0 && iblock <= block_size_i -1 + 0 && jblock >= 0 && jblock <= block_size_j -1 + 0) {
out[idx111] = (((gridtools::clang::float_type) -4.0 * lap[idx111]) + (__ldg(&(coeff[idx111])) * (((lap[idx111+1*1]
+ lap[idx111+1*-1]) + lap[idx111+stride_111_1*1]) + lap[idx111+stride_111_1*-1])));
 }
   // Slide kcaches

   // increment iterators
   idx111+=stride_111_2;
}}
```

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

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

Summary

- ESiWACE will extend the HPCW Benchmarks and apply them to various systems.
- ESiWACE made PSyclone and GridTools / DAWN work in/with production mode models.
- A full dry-dycore of ICON atmosphere was implemented within ESiWACE2. Early stage, but performance (wallclock time) is already promising.
- NEMO can be parsed by PSyclone.
- The interoperability demonstration of ESiWACE2 will allow to share backends/frontends being developed for different models.
- Automatic integration into Fortran allows to easily “escape” the DSL language for non performance critical operations or non DSL supported patterns -> allows to “port what you need”

ESiWACE2 has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 823988

ESiWACE2 is on Zenodo, the Open Access repository for scientific results
https://zenodo.org/communities/esiwace

Interested in getting in touch?
Twitter: https://twitter.com/esiwace
Website: www.esiwace.eu

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988