

# 2021-05-13

Yannick started the meeting with an update on JEDI 4. He announced that ewok had an update yesterday and should be able to run 4D Hofx now.

## JEDI 3

### JEDI 3 (Anna)

Benjamin has a series of PRs for BUMP major upgrade for the static B ([saber#85](#)).

Sergey extended observation-space localization test ([oops#1194](#)); added box-car and SOAR horizontal observation-space localizations that can be used with LETKF-type filters ([ufo#1132](#)); and a hybrid gain application ([oops#1207](#)).

Jo is working on a plan for the preconditioning changes for VarBC.

Anna is organizing a code sprint this summer, for documenting and cleaning up some of the oops C++ interface classes.

People who participate in the code sprint will pick one or two interface classes, and:

- add doxygen comments to the class and methods;
- check if the class implementation follows the [documentation](#), refactor it if it doesn't;
- check if the oops test for the interface class is covering all it should cover (and either create an issue, or amend the test if not);
- [optional] review L95 toy model implementation of the interface class;
- review other participants' PRs.

It is an opportunity to learn about some of the interfaces and oops code structure, as well as improve our code base! The work can be easily distributed.

Please let Anna know if you would like to participate.

## JEDI2

### JEDI 2 (Dan)

Jedi 2.1 - Interpolation

The JEDI1 team is looking at the unstructured interpolation and profiling it.

Jedi 2.3 - Use of NUOPC driver with FV3-JEDI

FV3 Cap to be refactored

Meetings taking place with ESMF team to design the refactor and gather requirements.

Jedi 2.6 - Ensemble DA validation

Initial data set was the wrong date. Jeff has prepared new data and provided the ensemble mean to act as the central for 3DEnVar.

Dan will prepare Yaml file with Aircraft for 3DEnVar.

SOCA team exercising the new LETKF, found some new bugs. Work underway to repair.

Jedi 2.7 - Background error model validation

Dan and Benjamin are looking into some issues with the psi/chi to wind transforms.

Jedi 2.8 - Regional DA

Prototype of hybrid DA experiment using JEDI and 13km. Fixed observations. Working to add interface to IODA in regional DA workflow. Workflow that runs GSI then generates IODA data and then uses observations in hybrid DA in JEDI.

#### Jedi 2.9 - MPAS general updates

Junmeio and JJ testing IODA V2 in the cycling experiments. Converted experimental data to IODA V2 format. Testing variational and h(x) applications for 120km 60 iteration. 5 AMSU and 5 conventional data types. Seeing significant slow down. 4 nodes 32 procs. Obspace constructor was 5/140 seconds, now 250/385 seconds. 60x increase in obs space constructor. Possibly to do with compression level 6 that is default in V1 to V2 converter.

#### Jedi 2.10 - LFRic general updates

No updates

#### Jedi 2.11 - UM general updates

No updates

#### Jedi 2.12 - Neptune general updates

Constructing a new Cylc-based workflow.

Cycling and debugging the cylc suites.

3DVar. Static B variances are too small in the tropics.

Cycled for about a week so far.

#### Jedi 2.14 - Cubed sphere grid into Atlas

Created a new base branch for the cubed sphere work.

Pushed latest develop into cubed-sphere-base.

Stripped out the UM-JEDI stuff, testing purely in Atlas.

Able to generate the GMesh figures.

#### Jedi 2.15 - VADER

Working through the prototypes. Close to having something working.

Adopted a Factory approach within the VADER repo.

#### Jedi 2.19 - Refactor GeoVaLs

Anna joined the meeting and reported that the allocation of GeoVaLs would be moved into the GeoVaLs class and that models would no longer perform the allocation.

Discussion with the model teams who all reported that the work should be minimal.

Met Office reported an issue where the GeoVaLs may have different dimensions for the increments. Marek and Anna discussed some work arounds that should be possible.

After Dan's update, Chris S reported that the MPAS team is now running regional DA. He said the implementation was very straightforward and commended the JEDI team for an extensible design. It was all done in configuration files with no code changes needed.

## JEDI1

The main focus of our meeting yesterday was two remaining development issues for the ioda v2 / JEDI 1.1 release. The first has to do with the format of ioda output files. In the context of diagnostics, [JJ and Junmei noticed](<https://github.com/JCSDA-internal/ioda/issues/250>) that some of the multi-dimensional variables in ioda output files were being written out in ioda v1 format (1D arrays) instead of ioda v2. This made us realize that, with all the focus on updating the `get\_db` interface, we had not properly updated the `put\_db` interface to handle the multi-dimensional ioda v2 format.

The long-term solution to this will involve a refactoring of the ObsVector class in ioda but we would not want to delay the release for this. Wojciech kindly offered to help with this by introducing a quick fix to `put\_db` that is similar to the one he previously implemented for `get\_db`, namely converting variables (excluding metadata) that end in `\_<number>` to 2D arrays before they are written out. This should be sufficient for this release since it covers all the current known use cases. It will allow us to put out a release where the input and output files are consistent.

The second development issue has to do with performance, but it is distinct from the ioda v2 performance issues that we have been discussing for the last few months, which have been resolved. This issue was separately identified by Mark O and JJ and manifests most clearly as a slowdown in the ObsSpace constructor for obs types such as Satwinds that have millions of variables. It has been traced to a problem with hdf5 in processing arrays of variable-length strings, such as datetime. Steve and Ryan are working on a solution and expect to have one soon.

We estimate that the release will occur about 3 weeks after these development issues have been resolved. Two weeks for documentation, clean-up & bugfixes after release branches have been created. One week code freeze.

Also regarding the release, Mark O has made a number of pull requests this week to upgrade mpas and other repos to be compatible with ecbuild 3.6+. Together with previous PRs by Mark and others, we're now nearly ready to make the transition to using the latest ecwf releases of ecbuild, eckit, fckit, and atlas in all containers and HPC modules. We plan to make these new containers and modules the default on all CI and HPC systems over the course of the next two weeks.

Maryam has been continuing to implement the publication of test results to cdash. She started with oops but now this has been implemented for most base JEDI repositories as well as the model repositories that are currently connected to CI, namely fv3-jedi and mpas. She is also working on archiving and documenting the tools she has developed for achieving this, including an AWS lambda function, by adding them to the jedi-tools repository.

Jacques is now up and running on Orion and ready to start running applications with ewok.

Summit/PGI update: fv3-bundle now compiles on Summit with PGI compilers and Mark M is now debugging some fv3-test failures. It appears that many test failures have to do with how PGI writes fortran namelists to the string arrays that are used by fms.

Azure update: the c192/c96 fv3-gfs 3denvar benchmark is now running on Azure clusters (without a container) and exhibiting comparable timings to Discover and AWS. The next step is to run a more computationally intensive application to better assess the performance of the different platforms. We're targeting JEDI-GDAS-10 for this but we want to wait until the ioda v2 PRs have been merged. Debugging of Intel OneAPI Supercontainers also continues.

Ben asked about Intel compilers in containers. Mark said that there have been ongoing discussions with NCAR/UCAR and Intel legal on this. The current status is that intel will allow the distribution of docker containers with Intel OneAPI compilers inside. Docker containers are multi-layered so it is clear which components came from intel and which did not. However, Singularity container images are flat binary files so this distinction cannot be made. So, when we update the intel containers for the release (likely next week) we will push a development container with compilers to Docker Hub. However, for Singularity, we make [a shell script and a dockerfile available on GitHub with instructions on how to use them to create your own intel development container](#). A warning however - it does take a lot of memory - at least 20 GB - to build the intel container. We will also continue to distribute [tagged containers to accompany each JEDI release](#).

Sergey asked if we have JEDI running on Azure clusters. Mark confirmed that yes, we do, on native modules - but not yet in multi-node singularity containers.

Ben then gave an update on **OBS 1** (for Hui). They are working on the JEDI GDAS application. They have modified routines to produce GeoVals files from JEDI and use them in the Hofx application. They are also working on calibration fixes. They are using the new ewok and seeing problems with the ozone - they are working on fixes. They are also working on speeding up diagnostics and the data ingest.

Francois (**OBS 2**) reported that the performance enhancement work continues with SageMaker. We have requested and obtained an increase in the number of user profiles for the Sagemaker instances, and thus the number of users that can run the diagnostics. Hailing is working on EDA at full resolution.

### OBS 3 (Ryan)

There were several OBS3 events recently.

- We held a general meeting yesterday. Notes are at <https://docs.google.com/document/d/1QJHZE8OpwgFSi07I61Byt9qsARS4bjzN1Fdn7qUfi8>.

There were two topics: First, Ron McLaren gave a talk on planned bufr library developments. Slides are linked in the meeting notes. In the remaining time, we discussed task plans and timelines for ioda-converters, ioda, and ufo.

- We started the ioda layout meeting series. Through Q2, we want to standardize what IODA data should look like. How are objects accessible to UFO and downstream users? What are the variable naming conventions? How do we want to use global and variable-specific attributes? This is a coordinated effort with many people from UCAR, EMC, and the Met Office, and we are discussing this again tomorrow.

- In conjunction with JEDI1, we sent out the upcoming release announcement for ioda v2. Mark has talked about the release already, but it's also worth mentioning that Steve has submitted six pull requests to various repositories for review. These range from the converters to the bundles. Once merged, our test data will be relocated to the ioda-data and ufo-data repositories.

For the converters, bufr, and odb:

- Cory Martin been upgrading the Python converters to write out the new file format version. His work focuses on the core Python library, so this will address all of our Python converters. Once we have new variable naming conventions, then we will update the dictionaries that map variables to their JEDI names, and this work can occur in parallel.

- Ron is extending the NCEP bufr library.

- Praveen has created a spreadsheet mapping the BUFR to IODA variables for all of the six ADPUPA bufr subsets. He is now working on creating a spreadsheet detailing the mnemonics from the ObsProc subroutine which reads all of the ADPUPA reports.

- David D. has made good progress on the ODB reader. Now that ioda-v2 is stabilizing, we are starting to port this code into the ioda-engines library.

For ufo:

- Michael Cooke has been working on expanding ObsFunctions to allow multiple where clauses to work correctly with the new variable assignment filter.

- Chris has been working on extending ObsSpaces in order to average profiles onto model levels.

- David S. has been expanding the UFO variable transformation code and is planning work on the adjustment of T / RH / wind at the surface with respect to model surface height. He is also finalizing a stationList to Yaml converter.

There was then some discussion of an OBS issue that Greg brought up having to do with the standard for the vertical coordinate: should we establish either a top-down or a bottom-up convention? Uncertainty can make it difficult to implement filters and other functionality. It was agreed that a standard is appropriate for GeoVals but the model grids will have their own custom implementation.

David S then asked about implementing a standard set of physical constants that can be used across JEDI repos. There was some agreement that this would be useful. Mark M wondered if udunits could provide this but it was unclear if the precision of the provided constants could be adapted as needed. It was agreed that the constants would have to be coded rather than specified in yaml files. We should discuss this further, potentially in a focus topic meeting.

Guillaume then gave a **SOCA** update. Travis is working on the Ewok workflow with SOCA. They ran into a science issue with GDAS reanalysis. Hamideh and Ming are updating to 2018 observations. Hamideh is also working on the sea ice Jacobian. They have LETKF running with the sea ice model. Kriti is working on regional ocean DA. Guillaume is working with the B Matrix, the ocean CICE model, and 3Denvar.



## LAND

### LAND (Andy)

#### FV3

Sergey, Clara and Tseganeh have ongoing work with the LETKF for snow depth, there's now a specific test for this

Need to add model orography/surface elevation to geometry so it can be used for observation localization and quality control, Dan has provided some guidance on how to do this.

Clara has found an issue with  $H(x)$  that is giving small, non-zero values for snow depth where we are expecting no snow. Assuming this is a feature of the interpolation rather than a bug, but we're checking that out.

#### WRF-Hydro NWM

Made sure everything is working with ecbuild3.6

Begun testing with IODA2, but don't think this will be a big deal for us

Mainly been trying to migrate to new(ish)  $H(x)$ , but this is requiring us to implement variable change, so taking a bit of time

We need to do this to handle the variable change associated with partitioning the multi-layer snowpack (which will also need to be done with FV3)

A lot of work has gone into understanding and designing this snowpack partitioning

#### UCLDAS

Zhichang has been working away on this, but not getting a big update on this until next Monday, so will report next time

### CRTM (Ben)

They are now finishing with the IASI coefficients which are missing short-wave (Patrick is working on this). They are also making comparisons between CRTM and RTTOV (Cheng is working on this). And, they are working on improvements to VIS and IR.