

2019-11-14

Yannick opened the meeting with no special announcements, and we proceeded with updates from everyone.

Boulder

Mark M has been working on an Intel compiler issue related to "hidden symbols". See [this JEDI team GitHub discussion](#) and [this ECMWF GitHub issue](#) for details. If you use the JEDI containers or environment modules then you do not need to worry about this - the issue is taken care of. If you have your own installation of JEDI dependencies, then Mark pointed out that the JEDI team discussion includes 3 options for working around this issue. Please select and implement only one of the options.

Mark M has a new environment module on Discover for Intel 17. Mark's module supplements those that Rahul has provided. Mark will help users get started for those interested.

Mark M has also been working with intel 19 environment modules on Cheyenne. The MPT stack installation failed because of build issues with HDF5 and intel. But, there is a working intel19-impi stack. ufo-bundle tests pass but the geos tests in fv3-bundle fail, apparently due to parallel io with netcdf4. Mark noted a curious symptom in that SABER tests using parallel HDF5 pass at the same time that FV3-GEOS tests fail.

Mark M reminded the group that a new Fortran OOPS variables module is available, deprecating the current Fortran module, and the model developers need to switch over to the new module soon. Mark noted that this has already been done for fv3-jedi and mpas. Marek reported that this has also been done for LFRic, and Travis reported that this has not been done yet for SOCA. See [this JEDI models GitHub discussion](#) for details and instructions.

Maryam is working on the ability to store and download large test data from Amazon S3 in our automatic testing.

Xin has a PR under review for variational bias correction, and another PR is coming soon. He is currently validating the automatic generation of bias predictors in the correction scheme.

Junmei has a PR under review which provides a workflow for running cycling experiments with MPAS using CYLC.

Mark O is refactoring jedi-rapids python source directory organization to accommodate the setuptools based methodology for building and installing python. Setuptools is a python provided packaging and distribution scheme that the group (core team and partners) have decided to pursue as our standard methodology for packaging and distributing JEDI components written in python.

Steve H encouraged the group to finish up work to bring all of the checked in netcdf observation files into compliance with the IODA conventions. There is a checker tool (tools/check_ioda_nc.py) in the IODA develop branch that gives you a PASS/FAIL message. Also, a ZenHub issue has been submitted for each file that fails the checker tool. If these have been assigned to the wrong person, please feel free to re-assign them to the proper person.

Steve H has a PR under review that enhances the ObsSpace Generate configuration which is used in conjunction with the MakeObs tool. This has resulted in a simple change to the format of the YAML file that will have to be done for anyone using the MakeObs tool in their tests. See [this JEDI models GitHub discussion](#) for details and examples of the new format. Contact Steve for help with this, or if you have any related questions.

Steve V is working on MPAS and related DA flows.

Travis added templates that assist the creation of an issue in GitHub. There are two templates added to the SOCA repository, one for enhancements and the other for bugs. When submitting a new issue, one of these templates can be selected which will automatically fill in the comment field with suggestions for appropriate content. The purpose of the templates is to help facilitate communication of helpful information to the person who will be working on resolving the issue.

Met Office

Marek reported that they have started work on the JEDI interface for NEMO.

Marek asked about the recently merged PRs ([UFO](#) and [OOPS](#)) concerning passing the MPI communicator to Fortran. It was noted that these PR's are particular to EDA, and updating now should not break current testing.

Marek mentioned that he will be updating the communicator in LFRic, and the GeoVals operator in the toy models next.

Marek introduced a new person at the Met Office, Wojciech migaj who will be working on organizing and processing observations.


David S reported that work on an obs operator for a new sonde instrument is underway.

EMC/GMAO

Ryan announced that he has a WCOSS account now.

Chris H is working on 3DVar and 4DVar for the shallow water model. The 3DVar capability has been merged in, and he recently fixed a bug in the adjoint interpolation which then got the 4DVar working. The 4DVar capability is in a PR which is ready to review now.

Discussion about CodeCov failures in the automated testing ensued. CodeCov check failures typically indicate that one has forgotten to add a test for the new functionality they just installed. In this sense, CodeCov is very valuable. We are, however, ironing out some kinks in the CodeCov reporting, so when you see a failure first check to see if it is indicating that you need to add a test (and please do so), otherwise check with the core team and see if anything needs to be done about the failure.

Guillaume noted that the templates for GitHub issues (reported earlier by Travis) came from NOAA. He also mentioned that they are working on reading the resulting H(x) values from FV3 into a SOCA DA run .

Hamideh is working on identifying the causes of differences between GSI and JEDI for $H(x)$ calculations. She has submitted ZenHub issues related to this work.

Hamideh reported that there exist some empty links on the data.jcsda.org website. Mark M has a tool to update links on the website. He will go run that tool after the meeting to get the empty links repaired.

NRL

Sarah announced that the Neptune grid has duplicate lat/lon values and they have been working to take this into account in the JEDI interface. They also reported back about our python discussion from last week. They confirmed that the visualization packages they use now are ready for python 3.6 and their cycl workflow implementation will soon be moving to python 3.6. So, as discussed last week, if we adopt python 3.6 as a required minimum version for JEDI tools and applications, then NRL will be ready for that.

Q/A

At this point the updates were complete and Yannick opened the floor to questions.

Travis reported that the GitHub issue templates don't show up in ZenHub, but he suggested that the templates should work in ZenHub. Mark M will look into whether ZenHub supports the GitHub issue template feature. If this looks compelling, we will try it out in the other JCSDA repositories. *(Update: Mark M and Phil confirmed that the templates are integrated with ZenHub and can be used when you open new issues in ZenHub; we will consider whether or not we wish to use such templates in the future but, in any case, it is good to know about them).*

Mark M announced that we are linking ZenHub issues into an AOP board through the ZenHub Epic mechanism. This is being done to improve our ability to report our progress to our stake holders. Mark reminded the group that when a new issue is created, to please fill out the Estimate and Epic (ie, link to an AOP Epic) fields need to be filled in. If the AOP Epics are not visible on your ZenHub board, you likely need to add AOP to the ZenHub workspace you are using. Please note that several workspaces (eg., DA, JEDI Observations) already exist that include AOP so it might be a matter of simply selecting one of those workspaces. Another tip is to add an AOP column to your ZenHub board which will help navigating through the collection of Epics and issues. Note - this primarily applies to JCSDA team members; in-kind and collaborative contributors need not worry about attaching issues to epics that concern the JCSDA Annual Operating Plan (AOP). In the future we may create new epics for in-kind and collaborative contributions and if and when we do so, we will let you know.

Marek asked for guidance with the ZenHub Estimate values. A non-linear scale, in units of "story points", is presented that you have to select from. One rule of thumb is to say that one story point is roughly a 1/2 day of work. Another tip is to think of the story point values as relative effort. For example the addition of a single ASSERT to flag something that was badly formed is a task that would be expected to take less effort than something like adding a new QC filter. In that case the ASSERT work should get a smaller story point value than the new QC filter. Regardless of how the story points are viewed, we are just starting out on this methodology and we will get better at estimating effort as we go along, and the key point is to get started now with the practice of providing estimates.

Ryan asked about doing a demo of profiling tools, since we still had about 20 minutes remaining in the meeting, and we decided to defer that to next Thursday's meeting (which is reserved for a special topic). This way other people that would like to demo can have a chance to prepare, as well as getting prepared with questions and discussion points about profiling.