

2020-05-14

Yannick opened the meeting by announcing the agenda: a general round-table update

He started by announcing that he has been working on a new series of pull requests that will re-organize the yaml files. **This will effect all models.** There was some discussion in the core team whether these yaml updates should be done incrementally or all at once. It was decided that the best approach may be to implement the changes all at once with coordinated pull requests, to avoid multiple disruptions of workflows down the line. Look for further details on this to come through the weekly meetings and through posts on the [GitHub JEDI Team page](#).

Rick then reported that, as posted by David H on the JCSDA slack page, AWS had a one-day workshop yesterday comprised of a series of 30-minute tutorials and webinars. Rick took particular interest in the new EC2 instance types they were describing and new storage options and resource management tools that should allow us to improve the cost effectiveness of our cloud usage. He is planning to investigate these new tools in more detail.

Mark M then mentioned a useful telecon that he and others (Tom, David, Francois) had with AWS representative Kevin Jorissen and Krthik Ramen. The main purpose of the call was how we might report on some of our joint work together but there was also some discussion about how we can improve the performance of multi-node JEDI applications on AWS. Motivated by this discussion, Mark, David, and Francois intend to implement AWS's new Elastic Fabric Adapter (EFA), which promises to deliver inter-node communication bandwidths comparable to infiniband. Mark intends to begin experimenting with EFA through AWS's ParallelCluster tool to see if this solves the multi-node performance issues we have been running into with a 3DVar benchmark. This will extend our current capability for launching clusters on AWS, which uses CloudFormation directly without ParallelCluster. David and Francois also plan to re-run the fv3-gfs forecast run we did last year with EFA.

Yannick mentioned that the refactoring of the yaml files he mentioned previously may help improve 3DVar efficiency by reducing the amount of memory needed through the elimination of redundancies. In fact, this is one of the motivations for re-organizing the yaml files.

Steve H then reported that he and Ryan are making good progress on their refactoring of the ioda ObsSpace and on providing new ioda interfaces. They expect to have a prototype ready for testing by the end of the month. Steve also mentioned that ioda-converters will use these new interfaces so some coordinated code changes will be needed to make the switch. They are considering scheduling a mini-code sprint for this. This is just a head-up for now but stay tuned for further details.

Yannick added that this new ioda layer has a python interface so ioda obs files can be read directly into python for plotting. This is a nice new feature that will be useful for diagnostics.

Nancy asked if these new python interfaces might be a good project for a student intern to work on. Ryan confirmed that they should be pretty straightforward to use for a student and added that he and Steve will provide some tutorials on how to use them.

Mark O then gave an update on the Near-Real-Time (NRT) website and JEDI-Rapids work. They received useful feedback from last week's presentation and are incorporating that feedback into the code. Mark is also looking into the issue of naming of buckets on S3. He initially avoided the use of dots in object names because this is not recommended for AWS S3. However, he learned that this recommendation does not apply to static websites so he is now considering using dots in the bucket names. So, be aware that bucket names may change in the near future.

Mark also reported progress on incorporating GNSSRO into the jedi-rapids/NRT workflow and expects to have that up relatively soon. He and the NRT team are also working to integrate FV3-GEOS into the NRT workflow. The goal is to use the same observation files for all models in the NRT workflow so that results can be compared. The observation files were constructed to properly align with the available GFS backgrounds for 3D variational analysis. The available GEOS backgrounds are offset (temporally) from the GFS backgrounds, so new (temporally aligned) GEOS backgrounds are required. The NRT team will work with Dan to provide the new GEOS backgrounds.

As Yannick asked for other issues, Mark M mentioned that his ufo-bundle build was failing this morning on his Mac. Steve noticed this too. After the meeting this was traced to how the long long int data type is defined on Mac operating systems. It is not a problem with linux. A fix is in place.

Then Iliana asked how issues like this are communicated - for example, when builds are broken. The first course of action is always to run a `make update` to get the latest changes. But, if there are still failures, look for a discussion on the [JEDI GitHub Team page](#). Chances are good someone else may have already posted the problem. If not, you can post it yourself. It is recommended that you "Watch" this page so you get email notifications whenever there is a post (see the drop-down menu on the left).

Guillaume recommended that JEDI consider a practice that they use in SOCA, which is to define stable branches that are slightly behind develop in order to minimize downstream build failures. Some discussion on this ensued. Mark O warned that this is more difficult for something like oops than soca because of all the downstream implications - not every workflow on every system can be tested. But Yannick was in favor of possibly implementing something like this after we set up automated application tests that build and test the model bundles when there are changes to base repos like oops, saber, ufo, or ioda. We're currently devising a strategy for implementing these integrated, automated application tests. Guillaume agreed that this functionality would help them out.

Yannick asked if there were any other updates, particularly from the model teams, but there were no responses. So the meeting was adjourned.