Mark and Yannick announced that there will be a re-organization of the JEDI1 and JEDI5 teams. The Infrastructure (JEDI 1) team will be split up into a Cloud/Workflow Infrastructure Team and a Software Infrastructure Team. Since they are tied to the AOP (which would require significant effort to change), we will keep the current ZenHub epics and labels. However, beginning in October, Epics JEDI 1.1-1.9 (cycling DA, NRT systems) will be part of the Cloud Infrastructure Team, led by Yannick. Meanwhile, JEDI 1.10-1.15 and 5.1-5.4 (testing, build issues, hpc modules, containers, release, documentation, academies) will be part of the Software Infrastructure team, led by Mark M. Some individuals who are currenly part of JEDI 1 will participate in both teams, others may wish to focus on one or the other. Yannick will initiate a new series of biweekly meetings for the Cloud/Workflow Infrastructure Team and we will sort out the participants in each team over the coming weeks. Yannick also mentioned that in the future, some of the NRT website work will be handled in a new JCSDA applications group that will lie outside of JEDI. But, that is further down the road.

We then discussed last week's testing code sprint and next week's upcoming documentation code sprint. Remaining issues from the testing code sprint that are not under review should be reclassified accordingly; high priority and in-progress issues should be added to the October milestone and the release. Medium-priority post-release issues should be assigned to the November milestone, and lower-priority to the Icebox.

We then reviewed the status of the milestones and discussed selected issues. C++14 support should now be available for all HPC modules and containers except for Hera. However, Ryan reported that NOAA has installed gcc 6.3 at our request so we can leverage that to provide C++-14 support for intel 18. That stack should be ready soon. Ryan also shared a list of comprehensive dependencies for JEDI and ufo that he compiled for NOAA and for WCOSS vetting in particular. this

There was a detailed discussion of several alternatives for providing python dependencies for JEDI. It was decided that the most promising approach is to handle the dependencies by repo, through setup.py scripts. Henry described how they do this currently with UFS. Repos that use CMake like ioda-converters can have CMake handle the package installation and linking to compiled binaries where necessary. But, the automated package installations should be optional; we should also provide a list of python dependencies in the jedi-stack README documentation specifying the python packages required by various components of JEDI. Then users will have the option to install these packages themselves based on their own preferred python environment (pip, conda, virtualenv, etc). On some HPC systems we may provide virtual environments through miniconda that users will have the option to use. But, this is not essential for the release. Ryan, Mark O, Rick, Henry, and Mark M agreed to work together to create focused issues and move forward with this plan.

Maryam had an update on CodeCov. It seems that it is free for public repositories so after our public release, we will be able to continue using it. So, there is no longer a need to consider other options.

Steve L reported that they have UFS building and running in a Singularity container using an extended version of the JEDI stack.

Jian is now running the NOAA global workflow s2s branch on Xserve/stampede2 (TACC) and Orion using the jedi-stack modules.

Action items identified at the close of the meeting include:
- Identify tasks for doc code sprint and add them to Epic 5.3
- Reclassify open issues from testing sprint (JEDI 1.15): milestone/release or icebox
- Move unfinished Sept milestone items to Oct or close them