

# Summary of the ESMF Change Review Board Meeting on August 31, 2012.

## Attendance:

Robert Ferraro/JPL, Cecelia DeLuca/NOAA/CIRES, Alan Wallcraft/NRL SSC, Tom Black/NOAA, Scott Sandgathe/NUOPC, Tim Campbell/NRL SSC/NUOPC, Atanas Trayanov/NASA GSFC, Gerhard Theurich/ESMF Core Team, Bob Oehmke/ESMF Core Team

## Agenda

The CRB covered the following topics during its meeting:

- Review development status

- Update development schedule

A synopsis of the discussion and decisions from the meeting is presented below. It is organized by agenda item. These notes attempt to capture the high points of the discussions, and any decisions that resulted.

## Development Status Review (Cecelia DeLuca)

The V5.3.0 release was announced on May 1. Three items expected for V5.3.0 were not complete at the time of the release, and have been deferred. V5.3.1, which was due in July, is not complete due to unanticipated time off for a couple of the key development staff.

Robert Ferraro noted that the Off—line Regridding capability has attracted a lot of attention in quarters outside of the normal ESMF user community. In particular, recent NASA data management proposals have identified it as a key technology to use for analysis tools and data set transformations.

The ESMF Core Team is proposing a change in the number scheme for ESMF releases. The proposed change is as follows:

- The 1<sup>st</sup> digit will align with Public Releases. Upon the delivery of a Public Release, the first digit will be advanced.
- The 2<sup>nd</sup> digit will align with Internal (Beta) Releases. This is a change from current practice.
- The 3<sup>rd</sup> digit will be reserved for patches to Internal Releases. The nominal value will be 0. If a patch to an internal release is required, the patched release will increment the 3<sup>rd</sup> digit by 1. No additional functionality will be included in a patched release.
- Public Releases will be designated by the highest numbered internal release upon which the Public Release is based, with a trailing “r” added. If the Public Release is patched, its numbering will be augmented by a patch number designator (e.g. “p1”, “p2”, etc.)

There was some discussion among the CRB members regarding the proposal, mostly regarding clarification of the proposed new scheme. It was generally agreed that having the major version number change in tandem with Public Releases made sense, along with reserving the last digit exclusively for patched versions of internal releases. Gerhard Theurich was given the action to write up a set of rules for release numbering, and send

them to the CRB for any final comment and approval. In anticipation of adoption of this new scheme, V5.3.1 has been renumbered V6.1.0. V5.3.0 should logically be numbered V6.0.0 under the new rules, but its current designation will be retained to avoid confusion. A “virtual” 6.0.0 may be linked to V5.3.0, but some further thought is needed to understand if this causes too much confusion among the user community.

Mesh subsetting was identified as a priority for the use of regridding by the data analysis community. It has been added as a new feature request.

Bit for Bit reproducibility is on the schedule for V6.1.0, and is a requirement for ESMF adoption by GFDL. It is also desired by the CRB for testing purposes. Currently, the sparse matrix multiply does runtime performance optimization that can result in unpredictable round off. Switching off the optimization should allow reproducibility, and is expected to be simple to implement if the number of processors is the same between runs. The more general problem (reproducibility across runs where the number of participating processors changes) is more difficult, since the code will need to enforce a specific order of execution. The V6.1.0 release will include reproducibility across runs on the same number of processors. The general capability will require additional study and development.

Tim Campbell asked where the regridding software was capable of identifying destination point which are unmapped in a regrid operation. Bob Oehmke said that the information was already available internally, and could be made available to the user. It was simply a matter of defining the interface for such a data return. Oehmke will contact Campbell to discuss the requirement. It is expected to be easy to implement, and has been tentatively included in the V6.1.0 release.

## **Current Schedule Review**

Release of V6.1.0 will be pushed back to Oct. The CRB decided to defer fleshing out V6.2.0 until the next meeting.

## **Next Meeting**

The next meeting will be scheduled to take place at the end of October, 2012, around the time of the V6.1.0 release.