

---

# George Mason University

ASSISTANT DEAN FOR RESEARCH  
CLIMATE DYNAMICS PROGRAM, MS 5C3  
SCHOOL OF COMPUTATIONAL SCIENCES  
FAIRFAX, VA 22030-4444

OFFICE: 301-595-7000

FAX: 301-595-9793

EMAIL: [schopf@scs.gmu.edu](mailto:schopf@scs.gmu.edu)

July 29, 2004

Dr. Tim Killeen  
ESMF Executive Committee  
National Center for Atmospheric Research  
1850 Table Mesa Drive  
Boulder, CO 80303

Dear Tim,

The ESMF Advisory Board has met following the highly successful ESMF Community Workshop and the ESMF Interagency Panel Meetings in Boulder. We would like to share with you our frank and considered opinion of the status of the ESMF project, some recommendations for activities through the final year of CAN, and our thoughts on the next iteration of ESMF. In short and above all, we are very impressed and pleased by the way things are going, but much work remains ahead.

## **Assessment of Project Status**

To assess the status of the ESMF project, we will examine the overall goals of the effort, comment on the scope of the project, then examine the degree to which the team is succeeding in meeting these goals.

### GOALS OF ESMF

The promise of ESMF is to

1. Facilitate the development of advanced models and components by providing a common modeling infrastructure that isolates developers from ever-changing computing architectures and provide shared repository of commonly used tools and functionality
2. Facilitate the coupling of models and components into more and more complex system models, improving interoperability and ease-of-use for system model evolution

These goals remain as valid and important now as they were at the project inception. In particular we feel that both items deserve attention, and that the Project's

approach to these goals seems appropriate.

## SCOPE

### *The Need to Improve as well as Couple Models:*

The goal to facilitate coupling receives quite a bit of attention, because it is always more exciting to do new things such as making novel combinations of models from existing parts, but this attitude can overlook the woeful state of our nation's current generation of models of the basic elements of the Earth system. Just coupling together more and more mediocre models and using multi-model ensembles to arrive at a consensus forecast has merit in a situation with a hard time-line and a pressing problem. But ESMF holds the promise of making it easier to actually improve the individual components, bringing the community together in more meaningful ways. The ESMF infrastructure has gone a long way to delivering a "Common Modeling Infrastructure" that was the goal of the group of scientists that started the ESMF process.

### *Superstructure Requires Infrastructure:*

In practice, the two goals of ESMF are quite closely related, since a key to model interoperability is the ability to exchange information between components using a well-understood yet flexible data exchange mechanism.

The design of the infrastructure (grids, arrays, fields, bundles and clocks) are obviously determined by the needs of the superstructure, and are essential to the successful execution of the coupling paradigm set forth in the ESMF\_Component classes.

What we appear to have at the current time is a very well-defined superstructure that has been demonstrated in the interoperability tests, has been found to work very well as a means to connect ESMF with the DOE CCA framework, and that serves to make the object orientation available at all levels of software development.

### *Grid Computing and Beyond...*

One of the issues raised at the Community Workshop was a question about whether ESMF delivers grid computing to the Earth science modeling community. It is the opinion of the ESMF Team, and we concur, that grid computing is not within the scope of the ESMF project.

That being said, ESMF seems well designed to facilitate grid computing solutions. Grid computing can be thought of as operating in two ways: first by providing a flexible pool of computing platforms that can be scanned for availability, allowing a task to be sent to one of a number of machines; second by allowing a job to be split across many machines, with the separate steps communicating across the network.

There are major software efforts underway to implement grid computing solutions in various forms, and it would be highly duplicative for ESMF to head off in

this direction. Yet, ESMF seems well designed for facilitating grid solutions. The flexibility of the virtual machine model in ESMF and the means for controlling domain decomposition on a single platform means that adapting a run to the local architecture should be easily implemented. If the second mode of operation is to come to pass, it is clear that the communication between tasks will need to be infrequent due to the high latency. What ESMF contributes to this mode of operation is a standard means for describing the data transfers and transforming coordinates between components.

#### PROGRESS, PROBLEMS, CAUTIONS

Having noted that the basic goals and scope of ESMF are sound, we see that there remain some significant challenges to the project.

##### *Limited Implementation of API:*

The infrastructure that is in place to make interoperability viable remains somewhat limited. We have previously encouraged the ESMF Project to concentrate on making at a version of the system that can work from end-to-end, and we are glad to see that this has happened. The demonstrations by the group at NASA Goddard using ESMF to “componentize” even the small functions of their GEOS GCM and to run using `esmf_2.0`, are very encouraging. In adopting this strategy, the team left some items incomplete, such as different grid types, vector re-gridding, and concurrent execution mode.

This choice has led to the meaningful demonstrations of ESMF, and we note the change in the audience attitude from the second to third community workshop: A year ago, the audience could be described a curious, skeptical, and technology-oriented. This year, the audience seemed more comprised of model users, wanting to learn how to make use of the system, where to get it, why it did not do this or that. The criticisms of ESMF at the workshop are in fact a great testimonial to the success of the effort so far: they reflect a strong desire to use ESMF and a frustration at having to wait for the full deployment.

##### *Responding to CAN vs. Responding to Users*

The CAN process that funds this cycle of ESMF places restrictions on the ability of the project to respond to the needs of users. The goals of the CAN process were admirable at the outset of a project like ESMF, because they ensured that a workable product would be produced, that the system would be demonstrated, and that the effort would not result in an academic exercise in software design.

As with all multi-year proposals, the team was able to set forth an accurate timeline and design plan for the first year, and able to do a respectable job identifying milestones for the second year, but it is difficult to see where one should be 4 years down the road.

Right now, the ESMF project has demonstrated interoperability. There are more interoperability experiments that are scheduled to be done, but it is fairly clear that not much will be gained from doing them. Nonetheless, the Project should

not risk failing to meet milestones. We recommend that an effort be made to engage NASA on a revision of these milestones. If unsuccessful, then the work will simply have to be done.

Users have expressed a desire for a concurrent execution mode, MPMD execution, more grids, additional IO capability. Others have asked for extensions to the ESMF framework to provide diagnostic toolsets, grid-computing enablement, integration with CCA.

We view ESMF as a framework, not a toolset. The ESMF Project has done an excellent job of scoping the effort by identifying the structure and the API. It should try to be responsive to groups asking that the API be implemented more fully, and come up with a prioritization methodology for answering these requests.

But we should recognize that a deployed ESMF core will provide the means for others to contribute and develop tools such as a diagnostics toolkit or to design a higher-level interface for grid computing. An examination of most software frameworks reveals that if the low level functionality is done well, others can build upon it to make more and more exotic and powerful applications. It is in this spirit that we caution ESMF to stick to the overall plan, have confidence in the extensibility of the product, and try to deliver the best functionality that serves the goals of the project.

#### *Open Source Practices*

We applaud the decision of the ESMF Project to embrace open source software for ESMF. This decision implies that others can easily build upon the core framework, or can even modify it for the benefit of the entire community. As has been demonstrated many times, a completely open code can still be kept very proprietary by simply changing the fundamentals very frequently. Outside contributors then find that they have to spend too much time modifying their work to keep up with new releases, and fall away from the effort.

We therefore encourage the ESMF Project to build as stable and robust an API and reference implementation as feasible, and to concentrate on the elements of ESMF that require high levels of communication between designers. As mentioned above, things such as diagnostic toolkits, GUI launchers, more grid definitions, grid enablement, etc. are prime candidates for the open-source community to contribute.

#### **Notes from the Community Workshop**

The Third ESMF Community Workshop was the most successful so far, and you are to be commended on the preparation and hard work that went into it. It is clear that ESMF is maturing at an excellent pace. Those who have tried out the system indicate that it is robust and working. A single day is really insufficient to convey the true nature and state of ESMF, and many questions had to be answered in a short-hand fashion. There is an excellent willingness on the part of the team to help the user community get started.

## **Highlights of Interagency Working Group**

The meeting of the Interagency Working Group was very helpful and encouraging. We are very encouraged to see initiatives proceeding in DoD, and see the possibility for strong links to NOAA, NSF and DOE. The discussion was aptly described by Don Anderson of NASA as “a first iteration” on the process of finding a suitable structure for ESMF follow-ons.

It is clear that a mechanism must be found for allowing the various agency interests to cooperate and coordinate for the furtherance of a unified ESMF. The Advisory Board would like to note the importance of the strong collaborative spirit behind the ESMF Project. If each agency funds a small set of ESMF investigators and developers, each of these sets will feel a need to distinguish themselves from the rest of the ESMF project and to demonstrate that what they are contributing is noteworthy and special. Such an attitude can easily lead to an unwillingness to fully share results, unwarranted criticism of other teams’ solutions, and a general disruption of coherence on the project. The agencies and the Project must work together at this point to ensure that the follow-on structure promotes coherence, and does not destroy it.

Solutions for this tendency do exist, such as the National Ocean Partnership Program (NOPP) which supports multi-agency work on HYCOM through a unified team. While they may not be applicable to ESMF, other opportunities may exist, such as the CCSP.

One valuable point that came out of the IAG meeting was that science must become a metric of how well ESMF is doing. The promise of ESMF is to enhance the scientific productivity of the Earth system modeling and data analysis community. In the end, ESMF will be judged in how effective it is in making this happen. There are competitors in the form of PRISM in Europe that may appear to deliver more rapid advances, so a challenge exists to deploy and highlight the contributions of ESMF to real-world science problems.

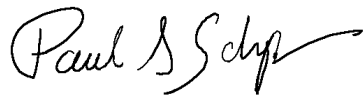
## **Summary**

As noted at the outset, we believe that ESMF is well on the way to success. You have developed a coherent concept and demonstrated that it can truly deliver the power that was dreamt of at the outset. Much work remains to be done during the last year of the CAN, but we are confident that the framework is clear, the work is relatively straightforward, and the team is talented. We recommend that you keep your efforts focused on the solidification of the API and the implementation of more of the well-established requirements. In areas where others can extend the framework to provide additional functionality, you should try to develop a mechanism to identify the needs and perhaps suggest a strategy for others to contribute.

Plans for an interagency-sponsored future are difficult to formulate in a vacuum. We feel that the strategy laid out at the interagency meeting is a good first draft of working arrangements, but perturbations will no doubt arise due to unantici-

pated needs and expectations of the various agencies.

Sincerely,

A handwritten signature in black ink that reads "Paul S. Schopf". The signature is written in a cursive style with a long, sweeping tail on the final letter.

Paul S. Schopf  
Chair  
ESMF Advisory Board