

---

# George Mason University

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

OFFICE: 301-902-1274

FAX: 301-595-9793

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

October 6, 2002

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

Dear Dr. Killeen,

The ESMF Advisory board would like to thank the Executive Committee, the Executive Board, and entire ESMF team for taking the time to outline and explain the goals, strategy and workings of the ESMF project. For many of us, this was our first close look at the operation, and without exception we were extremely impressed and encouraged. The team is as talented, experienced and capable a group of people as can be imagined for the effort. Your goals and vision are clear, and your approach to the problem of working as a team seem exemplary.

There is an emerging national policy thrust in climate science which aims to provide products for policy and decision makers. Of these, predictive models of the climate system are a clear and important tool. ESMF is a key mechanism for making the nation's effort in modeling more effective and productive.

We envision a future in which virtually all weather, climate, ocean and land models utilize the ESMF, both in the US and, hopefully, internationally, and it is this goal to which we commend your attention. We do not come to this view lightly, but rather through a careful consideration of what needs doing, what ESMF can deliver, and what alternatives could exist. The concept is so important that several centers started developing their own versions several years ago. Those talented people that undertook that earlier development have come together to form ESMF in a remarkable demonstration of cooperation for the common good. They have brought with them an incredible depth of experience and ability which convinces us of the likely success of the endeavor. The alternatives to ESMF are continued Balkanization of the Earth system modeling community, for we see no other group on the horizon ready to deliver on the promise of ESMF.

Our main message is carry on, keep up the good work and be assured that we are here to help, advocate and advise in whatever way we can.

Our more specific message is that the ultimate success of your endeavor will depend critically on new user acceptance, and we see challenges ahead as you enter

the software release stage. At one level, acceptance will come after the April 2003 software release, so you are urged to ensure that this target is met and in a way that is embraced by the community.

Nowhere does user acceptance come more to the fore than in the area of interoperability. We have the view that ESMF is comprised of two lofty goals: 1) provide a robust, portable, easily implemented toolkit which enables models to be built much more easily than they have in the past, and 2) use that toolkit to build model interoperability methods and standards which will encourage the sharing of model components. We are quite confident of your success in the the first within the initial 3 year period, however, the second is more difficult and will require thought, openness and good-will. The full expression of interoperability will take longer than 3 years.

You have asked for our opinion on a number of specific questions. We will try and give our replies below:

***Is the development process and software engineering approach appropriate?***

*If you are asking whether the utilization of software engineering in the first place is the right way to go,*

We feel that this project should only be carried forth in a software engineering approach. The scope of this effort should be to provide the tools, methods and infrastructure that allows model builders to focus on the physics and science of their problems. As a framework, it must be a software product built with solid engineering techniques. If you are successful, developers will rely on it for years to come, so it had better be done in a professional manner.

*If you are asking of our opinion of the type of software engineering approach that you have taken,*

We are extremely impressed by the development method as evidenced by the existing documentation, website and presentations. We are supportive of the heavily document-based approach and we feel that the time spent developing requirements, architecture plans and specifications will be well rewarded in the timely deployment of robust and testable codes and libraries. While lip-service is often given to planning first, then coding, it is satisfying to see the actual implementation of such a strategy and we are confident that this will prove the best, if not only way to carry this project forward.

***Is the technical strategy appropriate?***

You have raised some specific questions:

*What should be the role and extent of standards?*

We see that there are several levels of "standards" which come to the fore in the development of ESMF. At the lowest level, the use of standards-compliant compilers, communication libraries, and so on are essential to the success of ESMF. Current data standards, such as NetCDF, HDF, etc. should be supported, as the needs of the EVA suite and users dictate.

At the same time, there are developing (and often conflicting) standards for items such as variable names, flux directions, etc. The NetCDF CF standard, GRIB,

PRISM, Federal geospatial standards, among others, are being promulgated by outside groups. These standards typically provide names and/or codes for quantities, directions and units. We think that it would be a mistake for ESMF to engage in wholesale adoption of a single standard at this time, but that it should make provision for their use.

The board also sees a role for ESMF in proposing new standards where none exist, such as in developing grid descriptor conventions.

*Please comment on the architecture of an infrastructure/superstructure concept*

The strategy of focusing on both the infrastructure and superstructure aspects of ESMF is noted as a fundamental difference from the European PRISM project. We applaud this strategy.

The infrastructure provides the foundation from which the superstructure can be built, and a focus on the infrastructure is deemed essential to the success of ESMF. We also note that the infrastructure is based on the experience of FMS and GEMS and other similar efforts. Without an ESMF commitment to the infrastructure layer, continued development of separate infrastructure layers at several centers would go on, to the detriment and expense of the overall climate modeling community. It is especially encouraging that so many centers have agreed to abandon their own infrastructure efforts in favor of ESMF. On the other hand, the formation of ESMF has put continued development of these infrastructures on hold, so that it is imperative that this layer be delivered. Of course the developers of these systems are now forming the core of ESMF, so we are confident that the experience and intelligence needed to deploy a consistent and robust infrastructure layer lie well within the hands of the ESMF team.

At the same time, it is essential that the superstructure layer be a vital part of ESMF. The community will be looking to ESMF to lead the way toward new interoperability and model exchange scenarios that can be foreseen. While part of the infrastructure deals with the problem of isolating modelers from machine architecture issues, much of it is designed to facilitate the implementation of the superstructure. If we are going to alter the basic paradigm of how model experiments are carried out, it will be through the superstructure tools.

*Please comment on our approach to parallelization*

We see two important issues here:

1. Will the system be useful for the widest possible community of modelers?
2. Will the system provide true interoperability?

We encourage the ESMF to retain as broad an approach to parallelization as possible, supporting SPMD and MPMD models of execution, and multiple strategies for inter-processor communication. Barring strong technical reasons, the user should be free to implement either strategy for parallelism, depending on their machine architecture, problem mix and legacy code. We envision a future with national centers running high end climate models on thousands of processors, university and research centers running models on hundreds of commodity cluster machines, and individuals developing codes on desktop and laptop systems, all making fundamental use of ESMF infrastructure (in fact, this view is an accurate description

of the present use of GEMS, for instance). If too stringent rules are imposed within ESMF, you may end up with a system that is used at only a few centers, and which effectively divorces university and small center modeling from the national modeling strategy, and misses the opportunity of educating students who are ready to join research teams at the national centers.

On the other hand, we appreciate the value of model interoperability. We have had a glimpse of problems which may arise in regard to these issues. The Board would like have a report and/or discussion from the Executive Committee on these issues at its next meeting.

*Please comment on our decision on languages*

We understand the rationale for development in both F90 and C++ languages and feel that they represent a suitable basis for code development and user familiarity. The ESMF is urged to be aware of efforts to write high level couplers in alternative languages, such as python, although it is not clear at this time what direct action ESMF should take.

We further felt that insistence on F90 standards-compliance should be the design standard for Fortran ESMF codes, which is to say that ESMF should not go too far down the road to a "least common denominator" flavor of Fortran.

#### ***Monitor and Comment on Progress***

As noted above, we are very impressed with the current progress. We believe that the April 2003 release deadline will be met, albeit not without the concerted team effort, rapid resolution of internal debates and efficient coding practices that we know you to be capable of.

#### ***Monitor and Advise on relations to the outside community***

##### *NASA plugin program*

There is a clear need for extra funding to leverage the ESMF development process. This applies to NASA and OTHER agencies as well, though whether this should be earmarked funding or derived from existing programs is not clear.

##### *What is the role of the U.S. modeling and assimilation community?*

The input from individual researchers and independent teams will be highly beneficial to the overall goals of ESMF. We consider ESMF to be the first stage of the US Common Modeling Infrastructure development. As ESMF matures, the vision of the CMIWG will be able to be pursued with more clarity and focus. What is most needed is clarity in the ESMF April 2003 release, so that the rest of the community can see what you are about, can try out the software, and comment to you in a meaningful way.

While you are working feverishly to deploy the software, the community will likely take a much more leisurely pace towards evaluating it. For this reason, we see it important for the software release to be heavy on documentation. It should also include a Users Guide and tutorial. One of the more valuable aids would be templates of high level code. A strategy for reaching out to the community should be put in place after the April 2003 release.

### *PRISM, Frontiers and international relations*

The European effort PRISM has taken a substantially different approach to the problem of providing mix-and-match climate model interoperability. Their approach is designed to provide a rapid development path to interoperability, at the cost of generality (e.g. because of a premature choice of physical fields to be coupled). It has imposed structure and requirements on component models. In contrast, it seems that ESMF is trying to avoid these constraints, and provide a more open basis for future growth.

Our vision is of a situation in 5 - 10 years where we will be able to couple and exchange both European and American models. For this, PRISM currently looks like a first attempt to implement more general ESMF strategies. This is an acceptable situation, especially as ESMF will concentrate during the next 2 to 3 years mostly on the development of the so-called infrastructure layer and the general principles of the superstructure. This would allow then for discussion between ESMF and PRISM when ESMF will be ready to address the coupling superstructure layer and PRISM will be ready to start its second-phase development.

We do not see a need for a more formal relationship between ESMF and PRISM at this time. The groups should keep an eye on what each other is doing so as to avoid the future development of incompatibilities, which is readily obtained by cross-membership between the two groups.

We are not particularly familiar with the software engineering level of the Frontiers Project of Japan. We understand that NCAR will be attempting to run CSM on the Earth Simulator, and we anticipate that NCAR will do whatever work is needed in order to take advantage of the vector architecture of the Earth Simulator nodes. We also understand that message passing between nodes is consistent with the ESMF view, and that there is not a particular problem in this regard.

### *Vendors*

ESMF is engaged in developing flexible and robust software tools for the modeling community. This requires use of sophisticated high-level languages such as C++ and Fortran, and requires making extensive use of the newer features of these languages. Unfortunately, individual vendor implementations of these languages are not always as robust as needed. The solution to this impasse is the active involvement of compiler writers and engineers. How you get this involvement is not clear. As ESMF gets visibility and credibility, we anticipate that vendors will pay attention to the needs of ESMF and ensure that their compilers are standards-compliant. Much of this hinges on the size and visibility of ESMF.

### *Comment on Internal Collaboration*

We feel that the demonstrated success of the existing structure warrants no further change in the internal collaboration.

### *Discuss the Future of ESMF*

We feel that 3 years is a very short time in which to bring a project such as ESMF to a state of maturity, visibility and community acceptance that it should attain. It is imperative that some entity such as ESMF continue the work to maintain the infrastructure, entrain the larger community, and develop the new and exciting

functionality that perhaps can only be imagined at this time. What can be accomplished within 3 years is the construction of a solid foundation, a clear demonstration of its value, and generation of community support and awareness.

We feel that the larger community support will be essential for agency support of any follow-on effort. This is why we place such emphasis on delivering robust functionality on-time in the regularly scheduled software releases.

#### *Planning*

At the next Advisory Board meeting, we would like to engage in a discussion of what you see as a continuing level of support for any ESMF-like infrastructure, and how such infrastructure will (or will not) become community based.

#### *Next Meeting*

We would like to schedule the next meeting of the Advisory Board just after the April,2003 software release and public meeting.

#### *Closing*

In closing, we would like to congratulate the Executive Board on the excellent work that has been delivered to date. This has not been easy, nor guaranteed.

Since the Advisory Board meets only twice per year, feel free to contact me between meetings on any issue that might arise for which you may desire our input. I can easily contact the other members of the board and get back to you.

The advisory board would like thank the Executive Committee for their time and effort in briefing us. We would also like to give a kudo to acknowledge the special contribution, dedication, and leadership of Cecelia DeLuca. Her effort in facilitating communication and coordination and maintaining an extremely useful website have no doubt been of a great help to the effort. It was a great help to us.

Sincerely,

Paul S. Schopf, Chair  
ESMF Advisory Board