

# DRAFT

## Proposed CCSM – ESMF Evaluation Plan Stage1

*October 19, 2004*

### **1 Introduction**

This Evaluation Plan describes the first stage of the process that will be used to evaluate the suitability of incorporating the Earth System Modeling Framework (ESMF)<sup>1</sup> Version 2.0 or later into both the internal development and subsequent public distribution of the Community Climate System Model (CCSM)<sup>2</sup>. Incorporation of ESMF is defined as use of ESMF coupling software and select utilities.

In the first stage of evaluation, the ESMF architecture, ESMF redistribution and merging functionality and the ESMF time manager utility will be examined using the CAM standalone model. If this evaluation stage is deemed a success, then ESMF Version 2.0 or later will be officially incorporated into future CAM and CLM internal development (until the time at which a unified CCSM single-executable version removes the need for standalone CAM and CLM executables, when it will be incorporated into the CCSM).

The second stage of evaluation, which will be described in a later document, will evaluate the ability of ESMF to run the CCSM3 components CAM, POP, CLM, and CSIM, in either a sequential or concurrent single-executable (SPMD) implementation. This stage will evaluate ESMF regridding and concurrency functionality using a model system that is very close to CCSM3. Before the second stage of the evaluation is started, ESMF will demonstrate the ability to run gridded components concurrently and support the types of regridding and grids/decompositions used in the full CCSM model system. If the second stage of the Evaluation Plan is considered successful, ESMF will be formally adopted as a framework within which to construct flexible, sequential and concurrent single-executable versions of the CCSM system.

There are numerous motivating reasons why the creation of a single-executable version of the CCSM3 system utilizing the ESMF framework is desirable. In general, a SPMD version offers the following advantages:

- Scientific development across all model components (e.g. CAM, CLM, POP, CSIM) will occur in only one model system.
- Code maintenance and testing will be greatly simplified.
- SPMD is a “sweet spot” for vendor tools. The creation of a single-executable coupled model will simplify the process of error tracing and porting to new platforms.

---

<sup>1</sup> <http://www.esmf.ucar.edu>

<sup>2</sup> <http://www.cesm.ucar.edu>

Performance monitoring will also be improved since performance tools are normally designed to work with SPMD executables.

Reasons for evaluating ESMF to implement the single-executable version of CCSM include:

- The ESMF design is based on hierarchical components (e.g. CCSM contains CAM which contains physics and dynamics) and maps well to the CCSM. New science capabilities that are reasonably self-contained can be represented as new components and introduced in a more systematic way than they have been previously.
- ESMF provides a sophisticated online regridding capability that extends the regridding functionality currently in CCSM.
- ESMF offers a uniform communication standard that can be used among model components as well as within a given model component. This may help to clarify data exchanges within CCSM and reduce redundant code.
- ESMF should make it easier to add new fields to coupling exchanges.
- Together the items above allow for the straightforward implementation of potentially desirable features in CCSM, such as active chemistry and hierarchically nested models.
- Use of ESMF utilities such as the time manager, logging, and error handling can reduce code duplication and add new functions.

This Plan has been developed by members of the CCSM Software Engineering Group (CSEG) and the ESMF staff. CSEG will conduct the tests specified in the Evaluation Plan and will submit the final results to both the SSC and ESMF through the Software Engineering Working Group (SEWG), as specified in Section 4. If this Plan is approved by the SSC, then CCSM and ESMF staff will proceed with the evaluation.

## **2 Background**

CCSM currently uses an early version of ESMF, which we shall refer to as the ESMF prototype, for time management. The ESMF prototype has been bundled with the distributed version of the CCSM for about two years, starting with CCSM2. This prototype version has been and will continue to be supported by the ESMF team for the CCSM group.

ESMF development has continued beyond the ESMF prototype, and time management now represents a small piece of the framework functionality. Other utilities are included as well as coupling tools. Many of the interfaces in the ESMF prototype have changed. Extensive documentation and unit and system tests have been added to the ESMF distribution. The current ESMF distribution, Version 2.0, which we will refer to simply as ESMF, has many more files than the prototype and it is no longer feasible for it to be bundled with CCSM. The ESMF Version 2.0 software is currently being used in development and production codes at NASA Goddard, GFDL, MIT, CSU, NRL, NCEP, and other institutions. The NSF and DoD have committed to continued support and development of the framework, and additional funding is anticipated from NASA and NOAA.

### 3 Elements of the Evaluation

Four evaluation elements are defined in this document: the specific code that shall be used for the evaluation, the evaluation criteria, the performers of the evaluation, and the evaluation timeline.

#### 3.1 Code

The code will be a version of standalone CAM with modifications. The specific version will be a development model version and will be picked at the time the evaluation is to be completed. The ESMF component architecture will be used to structure the top, driver level of CAM, and the surface models used within CAM. The atmosphere and land (CLM) models will become ESMF components and will be called from an application driver level, and ESMF redistribution will be used to move data between the atmosphere and the land model components. Similarly, the CAM data ocean (CAM-DOCN) and CAM sea ice (CAM-CSIM) models will become ESMF components and will be called from the top driver level. CAM's Slab Ocean Model (CAM-SOM) will not be supported in this model version. If the evaluation is successful, CAM-SOM will be added in as an ESMF component afterwards. The prototype ESMF time manager will be replaced with the updated ESMF time manager. The evaluation code will contain CAM and CLM components that can still be run under CCSM. CAM will also be able to run the Single Column Atmosphere Model (SCAM).

The rationale for selecting the modified standalone CAM configuration to evaluate ESMF is as follows:

- It will be easier to test relative to using the full CCSM as a starting point for the ESMF evaluation. The standalone CAM with the same functionality, but without ESMF, will be used to evaluate correctness and performance.
- It is a reasonable starting point for constructing a single modeling application that would eventually have the ability either to be run as a sequential (i.e. standalone CAM) or concurrent implementation of the full CCSM. The CCSM project has an interest in developing such a configuration in order to reduce the costs of supporting both a standalone CAM and one which runs within CCSM.
- The standalone CAM has substantially more outside users and collaborators than CCSM. This will allow more people to evaluate whether the ESMF changes are viable.
- The insertion of an ESMF coupler between the CAM and the CLM is a first step in allowing the land model to be run on a different grid than the CAM. This is a long-term goal of some members of the CCSM project, and offers scientific incentive for introducing ESMF into the standalone CAM.
- Installing ESMF extensively within the fully coupled CCSM system would be a long-term and incremental effort, with many evaluations concerning appropriateness and correctness along the way. The cost of creating a fully coupled and validated

ESMF/CCSM system before performing an evaluation would be higher than performing the evaluation on this simpler standalone code.

The modified standalone CAM code described above will not test all the ESMF capabilities that may eventually be used, most notably regridding and concurrency. These will be addressed in the second stage of evaluation.

## **3.2 Criteria**

### **3.2.1 Platforms**

The CAM-ESMF evaluation code must support the platforms that CAM is supported and verified on. These platforms must be specified by CSEG by January 1, 2005. The CAM-ESMF test code must be verified to the same level that those platforms can be verified for the CAM version without ESMF.

### **3.2.2 Correctness**

The standard CAM test suite available at the time of the evaluation must show that answers only diverge by roundoff compared to the CAM code without ESMF. The CAM code with ESMF must also give bit-for-bit answers for differing numbers of tasks and threads. This must be shown on all CAM supported platforms and for all three CAM dynamical cores.

### **3.2.3 Memory**

Memory usage must not increase by more than 20% on a production IBM platform.

### **3.2.4 Performance**

We use the standard Eulerian dycore for all performance tests. Because we only use ESMF superstructure at model interfaces and not in the dycores themselves, we expect the overhead of using ESMF to be similar for all dycores.

The simulations shown in Table 1 will be run the number of times given in the table in order to correctly estimate the performance. ESMF must not increase time to solution of the evaluation code by more than 5% on the Bluesky IBM-SP and Phoenix Cray-X1 production platforms (for any of the configurations listed in Table-1). Performance will not be assessed on other platforms.

**Table-1. Matrix of Performance Assessment Tests to perform with and without ESMF on IBM and X1**

<i>Test</i>	<i>Resolution</i>	<i>#-sims</i>	<i>Tasks</i>	<i>Threads</i>	<i>Sim-length</i>	<i>Physics</i>
<b>1</b>	T42-L26	10	16	1	1 month	Default w 3 constituents
<b>2</b>	T42-L26	10	32	1	1 month	Default w 3 constituents
<b>3</b>	T42-L26	10	64	1	1 month	Default w 3 constituents
<b>4</b>	T42-L26	10	8	8	1 month	Default w 3 constituents
<b>5</b>	T85-L26	10	32	1	10 days	Default w 3 constituents
<b>6</b>	T85-L26	10	64	1	10 days	Default w 3 constituents
<b>7</b>	T85-L26	10	128	1	10 days	Default w 3 constituents
<b>8</b>	T85-L26	10	16	8	10 days	Default w 3 constituents

The rationale for selecting these tests follows.

Tests 1-4 are at the production CAM resolution of T42-L26. Tests 5-8 are at T85-L26 since this has been an important resolution for IPCC-type simulations.

We will evaluate scalability for the versions with and without ESMF, using a range of processor counts for each. We include a case with the maximum number of processors for each resolution as well as two scaled-down versions.

Tests 4 and 8 evaluate performance when using a hybrid MPI and OpenMP approach.

### **3.2.5 Interface Stability**

ESMF will not promise backwards compatibility of interfaces at this time. However, ESMF must guarantee sufficient stability for interface changes to not be disruptive to ongoing CCSM development.

### **3.2.6 Customer Support**

Prompt support from ESMF must be available to CCSM users and developers (response – but not necessarily fix - within 24-48 hours). Fixes must be available soon enough after being reported by CCSM staff to avoid disruption of or delays in the CCSM development activity.

### **3.2.7 Product Support**

ESMF must have resources for support and development of the framework for at least the next five years.

### **3.2.8 System Requirements**

ESMF must have the capability to be built with or without MPI.  
 ESMF must not require many other packages or exotic system software to compile or run.  
 ESMF must be straightforward to build on multiple platforms.

### **3.2.9 Software Requirements**

The ESMF team must satisfy software requirements identified by CCSM staff that are listed in the ESMF Requirements Document. These include the requirement that ESMF import and export state data types can use F90 pointers to reference the actual data used inside the model.

## **4 Staff and Process**

NASA ESMF funding for one FTE in CCSM at NCAR is guaranteed through this stage one evaluation. The evaluation code will be prepared by CSEG staff. The evaluation code will also be leveraged to satisfy NASA ESMF milestones.

CSEG staff will evaluate whether the criteria presented in Section 3 are met. The initial analysis will go to the ESMF team for a two-week response. This response time is intended to be sufficient to clear up any misunderstandings, but may not be sufficient to resolve outstanding issues. The results at the end of two weeks will be presented to CSEG and CCSM scientists. The ESMF team will have three months at that time to resolve outstanding issues. An evaluation report will be prepared at the end of that period that will be presented to the SSC through the CCSM SEWG.

Surveys on whether the modified standalone CAM code is aesthetically acceptable will be prepared for CCSM staff and external CAM/CCSM users. These will be reviewed by CSEG staff and scientists before being circulated. The results of the surveys will be presented to the SSC along with results from the CSEG evaluation.

Computational Science Section staff may be asked to help perform the set of performance evaluations.

## **5 Timeline**

The modified standalone CAM code is scheduled for completion by June 1.  
The evaluation and CSEG results will be available about 12 weeks later, by September 1.

## **6 Implications of Acceptance**

Stage one acceptance is defined as the ability of the evaluation code described in Section 3.1 to meet the criteria outlined in Section 3.2.

- Code that relies on ESMF will be incorporated into the main development branch of CAM and CLM.
- CCSM staff will begin the stage two evaluation. The performance, ease of use, portability, correctness, and other acceptance criteria of the ESMF software will continue to be monitored. The CCSM project reserves the right to decide not to use the ESMF software at a future time if some set of the criteria described in Section 3.1 ceases to be met and the ESMF project cannot provide a timely or adequate response.

- CAM and CLM users will need to download and install the ESMF software in order to compile and run future internal and public distributions of standalone CAM and CLM.

The rationale for the last bullet is as follows. It would be difficult to support two versions of CAM and CLM, one of which would be ESMF-free and another which would require the framework for compilation. Resources within the CCSM project are limited and supporting multiple versions of similar codes requires time and effort that may be better spent elsewhere.