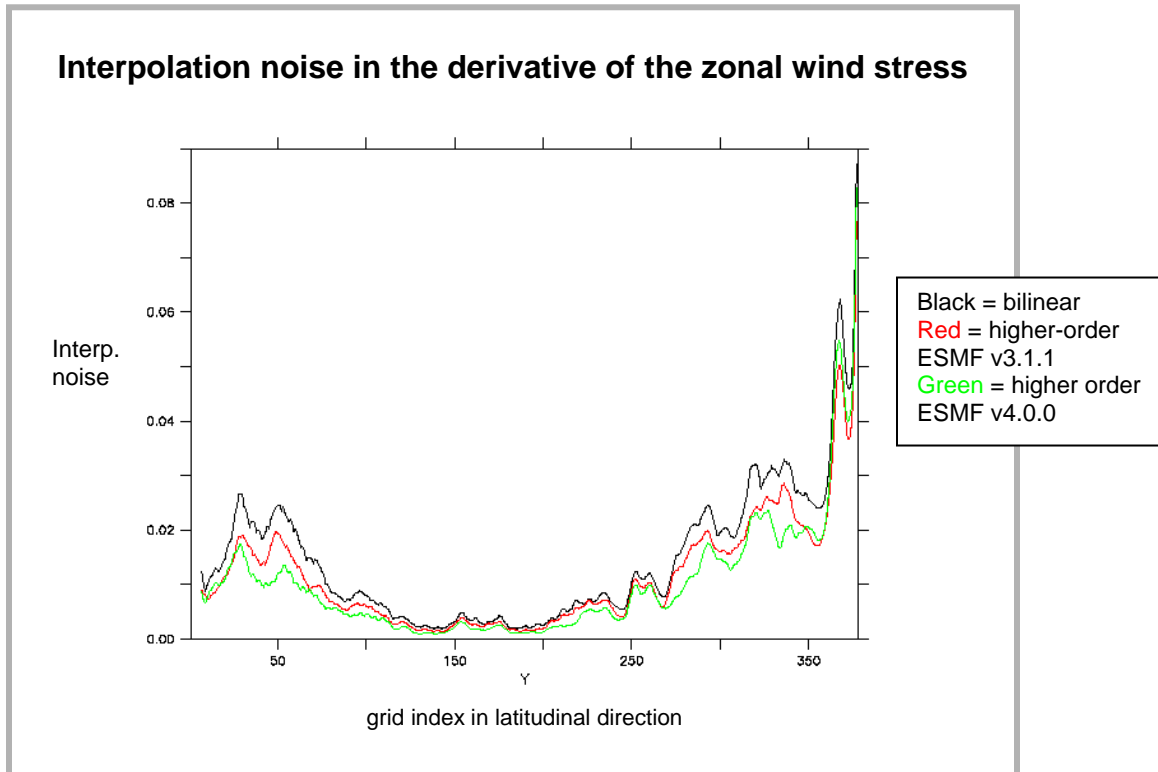


December 10, 2008



### ESMF Interpolation Technique Improves Representation of Ocean Transport in the Community Climate System Model

This graph shows the latitudinal structure of the interpolation noise in the derivative of the zonal wind stress ( $d\tau_{ux}/dy$ ). This derivative is closely related to the curl of the wind stress, which drives the upper ocean circulation and is therefore important to represent accurately in climate simulations. The wind stress was mapped from the T62 (2-degree) Community Atmospheric Model (CAM) grid to the gx1v5 (384x320 irregularly spaced) Parallel Ocean Program (POP) grid, using three different interpolation methods. The black line is the current CCSM bilinear method; the red line is an early version of the ESMF higher-order patch interpolation method (v3.1.1); and the green line is a recently improved version of the ESMF higher-order method (v4.0.0). The noise is calculated as the deviation of a point's value from the average of its own value plus the four surrounding points' values. The latest version of the ESMF higher-order method reduces the noise globally by 33% compared to the current CCSM approach. This is a significant improvement in the representation of transport processes in the model.

*(Image from Markus Jochum of NCAR Climate and Global Dynamics Division)*

## Table of Contents

1	Introduction.....	3
2	Software description .....	3
3	Timeline and strategic objectives .....	4
4	Accomplishments in brief for FY2008 .....	4
5	Plans in brief for FY2009.....	6
6	Detailed accomplishments for FY2008.....	6
6.1	Develop software for grid representation and regridding.....	6
6.2	Improve usability.....	7
6.3	Provide ESMF training and support .....	8
6.4	Develop software utilities and science gateway services .....	10
6.5	Framework optimization and porting .....	10
6.6	Advance community adoption .....	11
7	Project plan evaluation measures for FY2009 .....	12
7.1	Develop software for grid representation and regridding.....	12
7.2	Improve usability.....	13
7.3	Provide ESMF training and support .....	13
7.4	Develop software utilities and science gateway services .....	13
7.5	Framework optimization and porting .....	13
7.6	Advance community adoption .....	14
8	Impact of the ESMF project .....	14
9	Sponsors.....	14

## 1 Introduction

The Earth System Modeling Framework (ESMF) was motivated by the growing complexity of building and coupling Earth system models. ESMF provides a set of standard software interfaces and high-performance tools for common modeling functions, thereby promoting interoperable software systems and code reuse. Now in its sixth year, it has transitioned from NASA sponsorship to multi-agency support and management. It is the technical basis for the DoD Battlespace Environments Institute, the National Unified Operational Prediction Capability (NUOPC), the NASA Modeling Analysis and Prediction Program, and a host of smaller programs and projects. More information about ESMF and its applications is available on the ESMF website, <http://www.esmf.ucar.edu>.

ESMF has spawned a number of offshoot efforts, including the Earth System Curator project. Curator, which began under NSF funding and is now also NASA-supported, is focused on the development of web-based portals for assembling, running, and archiving climate model experiments. It emphasizes standardized metadata and semantic web technologies as a means of creating computational environments that are rich in stored knowledge and easy to use.

## 2 Software description

ESMF components are arranged in a hierarchical structure to form modeling applications. Exchanges are facilitated both for large, composite components, such as atmospheric models, and for smaller components such as specific physics parameterizations.

The current ESMF distribution contains:

- Tools for building scientific components and couplers, and a set of utilities for common modeling functions (e.g., calendar management, data communications)
- Concurrent or sequential execution, single or multiple executable modes
- Support for configuring ensemble members sequentially or concurrently
- Fortran interfaces and complete documentation, some C interfaces

ESMF supports 30+ platform/compiler combinations, including IBM, Cray X1, SGI, Linux, Mac, and other platforms. The complete list is available on the ESMF website.<sup>1</sup>

---

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

### 3 Timeline and strategic objectives

ESMF strategic objectives are described in detail in the *ESMF Strategic Plan for 2008-2010*<sup>2</sup>. As outlined in the *Plan*, there are currently three main activity areas:

1. Completing the functional requirements of the framework and improving its ease of use
2. Developing capabilities and strategic partnerships that help to integrate the framework into end-to-end computational environments (science gateways)
3. Supporting users who are integrating ESMF into applications

ESMF is now in its sixth year. Its second funding cycle under NOAA, NASA, DoD, and NSF sponsors will complete at the end of FY2010. In the next two years, the ESMF group anticipates two major release increments: ESMF 4, focusing on further rework and extensions to the grid representation and regridding software, and ESMF 5, focusing on standardization of interfaces and ease of use. Other deliverables in this time frame will likely include new and enhanced science gateways that utilize ESMF capabilities and components.

All ESMF activities are rooted in the project's core values of community ownership, distributed development, and open access to information.

### 4 Accomplishments in brief for FY2008

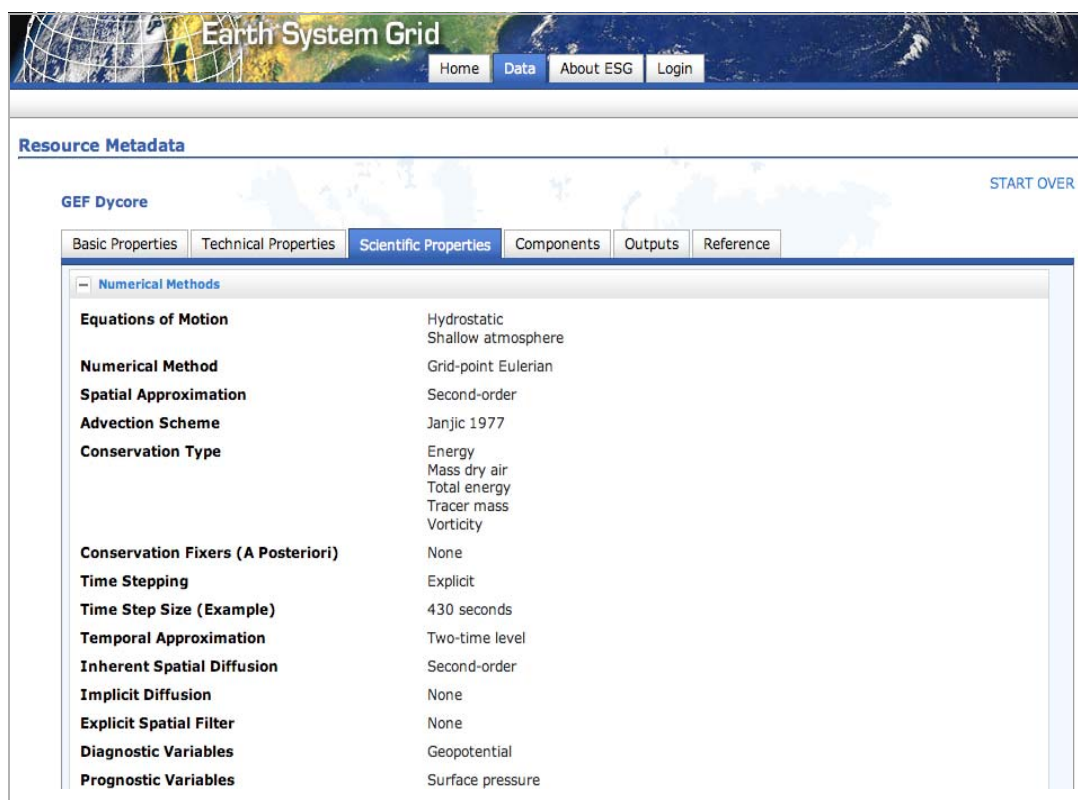
*Functional requirements and ease of use.* The ESMF team completed the framework's transition to a new set of central data structures in FY2008. The aim of this redesign effort, which began in 2005, was support for a wider variety of supported grid types and general, parallel regridding. There were two major releases. The first, ESMF 3.1.0r (May 2008), was a robust and portable public release that enabled production codes at NCEP, NRL, and NASA to transition to a basic version of the new data structures. The second, internal (beta) release ESMF v3.1.1 (September 2008), contained two important new grid types – unstructured meshes and observational data streams. This release also introduced the capability to generate bilinear and higher-order interpolation weights in parallel, in up to three dimensions. The higher-order interpolation strategy is being evaluated by the oceanography section at NCAR and has shown good preliminary results in reducing interpolation noise in key variables (see cover figure).

*Science Gateways.* The Curator and ESMF teams partnered with the DOE-funded Earth System Grid (ESG) project and its NSF-sponsored affiliate, the Science Gateway Framework (SGF), to develop a prototype portal for the 2008 NCAR Advanced Study Program Colloquium entitled *Numerical Techniques for*

---

<sup>2</sup> [http://www.esmf.ucar.edu/plans/plan\\_0802\\_esmfstrat.doc](http://www.esmf.ucar.edu/plans/plan_0802_esmfstrat.doc)

*Global Atmospheric Models.* Over two weeks, about 80 Colloquium participants ran 13 atmospheric dynamical cores – the portion of an atmospheric model that solves the fluid equations - in specified test configurations, and compared their output. Workshop organizers and technical collaborators extended the ESG ontology with additional metadata describing each dynamical core in detail. They used the extended ontology to create an on-line environment in which the hundreds of collected datasets could be analyzed collaboratively. Capabilities included tools and user interfaces for faceted search and browse (from ESG); plus links from datasets to exhaustive descriptions of the components that generated them and dynamic comparison tables (developed by Curator). Figure 1 shows a trackback page – the component description associated with a particular dataset – displayed in the portal.



**Figure 1. Dataset trackback.** Detailed description of one of the dynamical cores used in a 2008 Summer Colloquium at NCAR, which is linked to experiments and datasets. This webpage is from a prototype portal developed collaboratively by partners including the Curator and Earth System Grid projects.

*ESMF applications.* The number of ESMF science components in the community is an important metric, since more compliant components mean more options for researchers creating coupled systems. The adoption of ESMF grew steadily this year, with the number of available science components growing from 58 at end

FY2007 to 69 at end FY2008. A full listing of components is available on the ESMF website.<sup>3</sup>

## 5 Plans in brief for FY2009

*Functional requirements and ease of use.* A primary goal for FY2009 is delivery of a public release (ESMF v4.0), that includes meshes, observational data streams, and online regridding. These features, which were delivered in beta release v3.1.1 during FY2008, require additional testing, optimization, and standardization. With this public release, the ESMF development team will have satisfied most of the project's initial functional requirements. The development focus then will be on completing the ESMF package and preparing it for extensive use, including deployment at operational centers throughout the U.S. as part of the NUOPC program.

*Science Gateways.* ESMF plans to continue collaborations with the Earth System Grid, Earth System Curator, NASA MAP effort, and other programs and projects implementing workflow and knowledge systems. GFDL is also a key partner, for the tools and metadata it has developed related to grids, and for its continued innovation and leadership in the development of end-to-end modeling environments.

*ESMF applications.* The ESMF team anticipates completing the ESMF-compliant version of CCSM in FY2009. Key features to be added in support of this and other applications are conservative regridding options and parallel IO.

## 6 Detailed accomplishments for FY2008

During FY2008, the ESMF team made significant advances in software development; training and support; and community adoption. The following sections describe metrics of success and accomplishments in detail.

### 6.1 Develop software for grid representation and regridding

#### Measures of success:

- Delivery of ESMF releases capable of representing multi-patch logically rectangular grids, unstructured grids, and observational data streams.
- Delivery of an ESMF release that generates regridding weights internally and can perform regridding using those weights.

#### Actual performance, outputs and outcomes:

Internal (beta) release ESMF v3.1.1, completed in September, 2008, contains a new class that represents unstructured Meshes. In this release, the framework

---

<sup>3</sup> <http://www.esmf.ucar.edu/components/>

can generate bilinear and higher order weights in parallel, in up to three dimensions. The release also includes a Location Stream class (contributed by NASA GSFC) for the representation and manipulation of observational data. Meshes, Location Streams, and (logically rectangular) Grids have been implemented so that a Field or a FieldBundle can associate data with any of these underlying topologies using a similar interface. Another capability added this year was the ability to perform direct coupling between components; coupling can now happen without needing to return control to a coupler. This was desired by several groups for use cases ranging from I/O to ocean biogeochemistry. Field and Bundle communication methods that were disabled during data structure rework were restored in v3.1.1. Although beta v3.1.1 is the first release to contain many long-awaited capabilities, the public release prior to it, v3.1.0r, was also an important milestone. The v3.1.0r release, in May 2008, was an extremely stable and portable package which has enabled many ESMF customers to transition to the new data structures.

A physical-space representation of multi-tile grids is still forthcoming. However, multi-tile grids can be represented in index-space in recent releases. Additional testing, consistency checks and other preparation for public release will be required in the coming year for Mesh and Location Stream classes, and Mesh classes will need C interfaces in addition to Fortran in order to satisfy customer requirements. Implementation of conservative regridding methods will also be needed.

## **6.2 Improve usability**

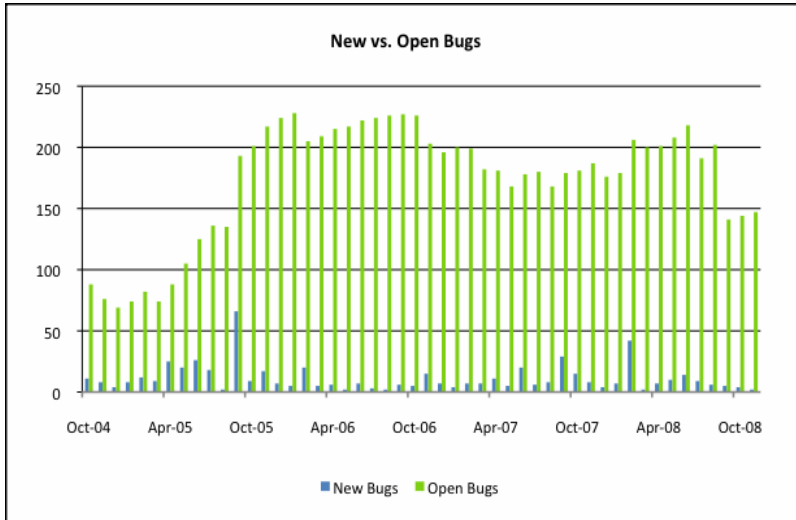
### **Measures of success:**

- Reduction by half of outstanding bug reports.
- Delivery of a release that further improves consistency of behavior and interfaces.
- Delivery of C/C++ interfaces.
- 100% test coverage of framework methods.

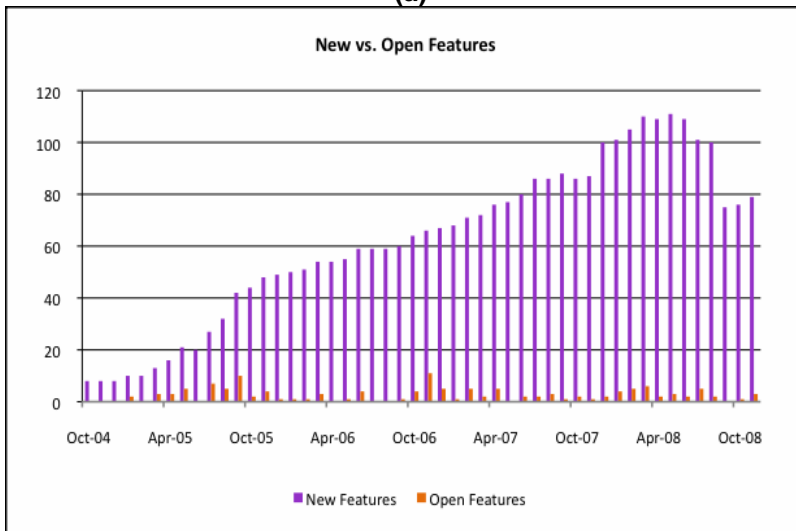
### **Actual performance, outputs, and outcomes:**

The latest beta release of the software (v3.1.1) is at 96% test coverage, with several of the new Mesh and Location Stream methods not fully tested in addition to a few other stragglers. These will be fully tested before the next public release. As shown in Figure 2 (a), the number of open bug reports was reduced by 40% in September, 2008. Figure 2 (b) shows that in the last quarter of FY2008, the number of open feature requests decreased substantially for the first time in the project's history, by about 30%. This was due to the many capabilities finally released with ESMF v3.1.1. Releases this year included standardization of inter-language interfaces and logical data types. C interfaces were completed to support an initial coupling between pWASH123 and ADCIRC, in which the interpolation weights are computed outside ESMF.

Standardization of error handling was deferred to next year. C interfaces for the new Mesh class will be developed in the future to support a coupling of pWASH123 and ADCIRC in which weights are calculated by ESMF.



(a)



(b)

**Figure 2. Bug and feature request metrics.** Chart (a) shows a 40% reduction in open ESMF bug reports from peak. Chart (b) shows a 30% reduction from peak in ESMF feature requests. FY2008 was the first year that the number of ESMF open feature requests declined, due in large part to the new grid and regridding capabilities that were delivered.

### 6.3 Provide ESMF training and support

**Measures of success:**

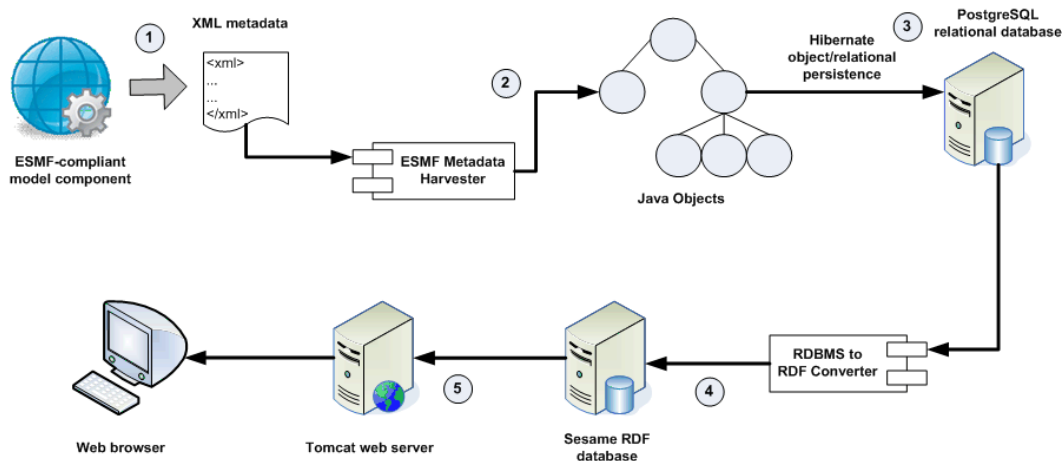
- Conduct two ESMF training classes.
- Bring ESMF documentation up to date with new grid development.

- Develop system tests and/or demos showing a variety of configurations of ESMF applications, including ensembles and alternative methods of time management and sequencing.
- Submit a journal or conference paper on the ESMF design or processes.
- User comments, surveys, and/or evaluations show that ESMF customers are satisfied with support, training, and documentation.

### Actual performance, outputs, and outcomes:

A hands-on, multi-day “coding workshop” tutorial session was part of the 7<sup>th</sup> ESMF Annual Meeting, held coincident with the AGU Joint Assembly in Fort Lauderdale, FL on May 28-30, 2008. This enabled a small group (~12) of customers to work intensively with ESMF developers, mainly on the transition to new data structures. The tutorial focused on the transition to new ESMF data structures. Feedback collected informally was positive. A second tutorial was deferred in order to focus on finishing outstanding functionality.

ESMF documentation was brought up to date with new grid, unstructured mesh, and observational data stream development. Many improvements were made to documentation based on support requests and other customer feedback. Examples were developed for new grid and regridding capabilities. A number of system tests were added, which included demonstrations of ensemble support, recursive components, and coupling of multiple executables. A journal paper on ESMF processes and governance is in progress. A journal paper on the Earth System Curator project was published.<sup>4</sup>



**Figure 3. Self-describing models.** This figure shows how descriptive metadata written out by ESMF is ingested into a prototype web portal and transformed into a semantically useful format before being displayed at the portal’s user interface. (Image from Rocky Dunlap, Georgia Institute of Technology)

<sup>4</sup> Dunlap, Rocky, L. Mark, S. Rugaber, V. Balaji, J. Chastang, L. Cinquini, C. DeLuca, D. Middleton and S. Murphy, “Earth system curator: metadata infrastructure for climate modeling,” *Earth Science Informatics*, Volume 1, Numbers 3-4, November, 2008.

## 6.4 Develop software utilities and science gateway services

### Measure of success:

- Delivery of a public ESMF release that includes new features for existing utilities, prioritized in response to user input.

### Actual performance, outputs, and outcomes:

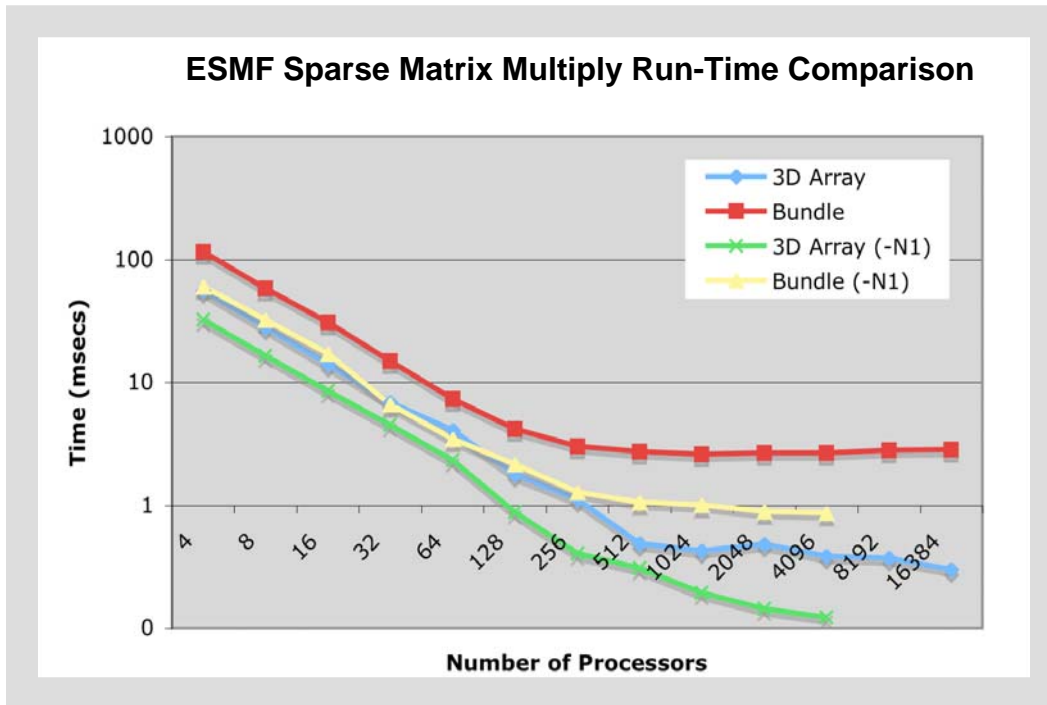
In the Time Manager, a modified Julian date calendar was implemented at the request of space weather groups, and several bugs were fixed.

A productive student intern completely reworked the Attribute class, which is used to associate descriptors in the form of name/value pairs with a variety of different classes. In the v3.1.1 release, Attributes are newly hierarchical (e.g. the Attributes of a FieldBundle aggregate all the Attributes of the Fields it contains). ESMF now supports standard Attribute packages, which are sets of Attributes that are predefined for a particular purpose and following a stated convention. This was done in order to improve model documentation and interoperability. Attributes can be written out either in plain text or in XML. The XML version can be used to link model descriptions with datasets in a prototype distribution portal developed in collaboration with the Curator, ESG, and SGF projects. Together these capabilities are an important step towards self-describing models. Figure 3 shows the transformations that occur when ingesting ESMF-produced XML metadata into a portal, populating a relational database with it, and then harvesting the metadata into RDF (Resource Description Framework) form for use with tools that create search, browse, and other user interfaces. RDF is a format upon which reasoning capabilities can be built.

## 6.5 Framework optimization and porting

### Measures of success:

- Minimal performance burden for working code. (Target <5% overhead in component overhead and regridding methods.)
- Benchmark regridding methods for a range of processor counts (1,000-5,000 processors) and show good performance and scalability.



**Figure 4. Performance at the petascale.** The chart above shows scaling of the ESMF sparse matrix multiply, used in regridding transformations, out to 16K processors for two ESMF data structures. The ESMF 3D Array scales out to 8K processors and is the recommended strategy for petascale computations. -N1 means 1 core per node. (Image courtesy of Peggy Li, NASA/JPL)

### Actual performance, outputs, and outcomes:

The Array sparse matrix multiply method was benchmarked out to 16K processors. Two data structures were tested: an ESMF 3D Array, in which the data from a set of 2D Fields was stored contiguously in a single native array, and an ESMF ArrayBundle, in which the data from each 2D Field is stored in a separate native array. The 3D Array showed good scaling to 8K processors. The Array Bundle will be subject to further optimization. See Figure 4 for the scaling curves.

A performance report was posted in July 2008 showing substantial memory improvements from recent optimizations.<sup>5</sup> Notable new ports this year included gfortran and Solaris.

## 6.6 Advance community adoption

### Means, strategies, and actions:

Assist in the development of new, coupled ESMF applications and the transition of these applications to routine use and operations.

<sup>5</sup> [http://www.esmf.ucar.edu/metrics/performance/timing\\_0807\\_memory\\_scalability.pdf](http://www.esmf.ucar.edu/metrics/performance/timing_0807_memory_scalability.pdf)

### **Measures of success:**

- Number of ESMF components available.
- Number of ESMF applications in routine use or operations.

### **Actual performance, outputs, and outcomes:**

There are now more than 69 ESMF components available, a significant increase over the 58 available at the end of FY2007.

Application development In FY2008 focused on incorporating new ESMF capabilities into codes rather than transitioning codes to operations. We expect FY2009 to be marked by the transition of more ESMF codes into production and operational use. We are aware of four major codes that have put ESMF into extensive and routine production use. These each include multiple components:

- The NASA GEOS-5 atmospheric general circulation model based on ESMF has been in production since FY2005.
- The NCEP Global Forecasting System using ESMF was put into operations on August 22, 2006.
- The coupled HYCOM and CICE model using ESMF is now in routine use at the Naval Research Laboratory.
- The COAMPS and NCOM application using ESMF components and coupling has been in use since FY2007.

The National Unified Operational Prediction Capability (NUOPC) is one of the larger programs to be based on ESMF. It is intended to enable a Tri-Agency (NOAA, Air Force, Navy) global atmospheric ensemble system. It got underway in FY2008 with the formation of committees who will recommend strategies for system development. The ESMF Core Team manager (DeLuca) is a member of the Common Modeling Architecture Committee, which is developing recommendations for coding tools and standards. This activity is likely to generate new ESMF components and applications in the coming years.

## **7 Project plan evaluation measures for FY2009**

The sections below describe key metrics and goals for FY2009.

### **7.1 Develop software for grid representation and regridding**

#### **Measures of success:**

- Delivery of an ESMF public release that includes the grid and regridding capabilities released in ESMF beta version 3.1.1 in FY2008.
- Delivery of an ESMF release with a high-level interface for the representation of multi-tile grids and exchange grids.
- Delivery of an ESMF release that offers on-line conservative regridding.

## **7.2 Improve usability**

### **Measures of success:**

- Delivery of a release that further improves consistency of behavior and interfaces.
- Complete sufficient C interfaces to meet customer needs.
- Maintain or reduce the number of open bug reports and feature requests.

## **7.3 Provide ESMF training and support**

### **Measures of success:**

- Hold two training classes during FY2009.
- User comments, surveys, and/or evaluations show that ESMF customers are satisfied with support, training, and documentation.
- Produce a journal paper on ESMF.
- Develop an Outreach Plan to promote use of the framework.

## **7.4 Develop software utilities and science gateway services**

### **Measures of success:**

- Delivery of a public ESMF release that includes new features for existing utilities.
- Delivery of a basic I/O library.
- Continued development of ESMF Attributes, including a more general and robust XML write capability.
- Development of web service bindings for ESMF standard interfaces. This will enable components to be integrated more easily into heterogeneous environments that include model execution, visualization, analysis, etc.

## **7.5 Framework optimization and porting**

### **Measures of success:**

- Minimal performance burden for working code. (Target <5% overhead for ESMF component and regridding methods vs. native packages.)
- Continued demonstration of performance viability for petascale and other emerging computing architectures.
- Requested ports delivered, including Windows.

## **7.6 Advance community adoption**

### **Measures of success:**

- Number of ESMF components available. A goal for FY2009 will be completion of an ESMF version of CCSM4.
- Number of ESMF applications in routine use or operations. Goals for FY2009 include having at least two more ESMF codes transition to production or operations.
- Number of programs based on ESMF.

## **8 Impact of the ESMF project**

Widespread use of ESMF represents a paradigm shift in the way weather and climate models are constructed. Through increased code interoperability, community building, and standard tools, ESMF is beginning to make model development easier and to facilitate new, multi-agency science collaborations. The end result is an Earth science community better equipped to explore basic research issues and better equipped to answer questions about the impacts of Earth science on society.

## **9 Sponsors**

Core ESMF development is sponsored by the National Science Foundation, the National Aeronautics and Space Agency, the National Oceanic and Atmospheric Administration, and the Department of Defense. In addition, a variety of ESMF-based application adoption projects have been sponsored by NASA, NOAA, the U.S. Geological Survey/Department of the Interior, and other agencies and institutions.