EXECUTIVE SUMMARY

This EarthCube Roadmap project explores new opportunities for adapting, relaxing and enhancing various model coupling frameworks used in Earth System Models (ESMs) in order to accelerate broader utilization of these models by the rest of the Earth sciences and decision making communities. We advocate here for a 'Geoscience Community Modeling Framework' to accelerate progress on a multitude of aspects related to modeling. We advocate building upon and integrating several of the modeling frameworks that already exist while accommodating new frameworks where necessary, and stress the need for cooperation between the various modeling framework projects in order to identify and adopt (or mediate between) common standards for dealing with a variety of interoperability issues including: (1) model metadata, (2) semantic mediation and (3) formats for describing unstructured grids. The core of this Roadmap is a proposed 2-solution strategy that will:

1. **Develop and integrate a set of standards-based protocols for linking component models** of a myriad of Earth System processes into a coherent, conservative, multi-scale and high-performance computing enabled modeling system, and
2. **Develop an Earth System Model coupling testbed** to facilitate standards-based approaches to ESM component coupling and to support broader engagement and participation of geoscience domain scientists in community-based ESM development and end-use applications.

In proposing this solution we were guided by several overarching principles which sought to a) optimize engagement and participation across geoscience and relevant end-user communities, b) capitalize on existing and sometimes long-standing model development efforts, c) utilize standards-based approaches to coding and semantic ontologies and d) fully engage current and future computational capabilities.

Given the technical and institutional challenges that exist, it is realistic to consider a solution for Earth System Model coupling which is adaptable and evolutionary in the sense that the solution can change over time in response to the demands of the community relying on it rather than presuming to design a singular end-to-end solution that will last in perpetuity.

The roadmap articulates a set of requirements or ideal attributes for developing new modeling frameworks which will facilitate efficiencies in Earth System component model coupling. One principal goal of this modeling Roadmap is to advocate for the concept of 'community-based' modeling that supports open model development processes, and community-defined modeling use cases and assessment benchmarks.

In producing this Roadmap, we reviewed the status and capabilities of several current model frameworks, with inputs from their developers and users, and develop a process by which different model coupling activities can evaluated. One focal point of our 'process' is to develop a simple worksheet or survey which will be completed as the coupling project proceeds. We seek to develop standardized questions to obtain input which would be utilizeable to the immediate activity as well as a broader suite of coupling activities. This worksheet will be drafted and evolved through the execution of a set of pilot studies or 'use cases' in model coupling, first focusing on pilot projects explicitly designed to examine issues related to the coupling of terrestrial hydrology models with other models of the Earth System. We propose a basic timeline of initial assessment activities that will deliver an

articulation of a suite of model testbed attributes and metrics that are most suitable for Earth System Model coupling applications. Initially we propose that testbed archive and maintain inventories of model and coupling framework codes and their documentation, datasets for a growing database of use cases and archives of previous coupling exercise experiences as documented by our assessment worksheets. Once the process is finished, a basic documentation and evaluation framework for assessing model coupling activities would also be in place and, if designed correctly, the framework will have an evolutionary capability which allows it to change with the changing requirements of ESMs.

1. ***Purpose:*** *Introduction, including community(ies) to be served, technical area(s) of the roadmap, and brief discussion what improvements in the present state-of-the-art in the proposed capability it will enable. Also include examples of how the outcomes from your effort will enable the community to be more productive and capable.*

   The concept of an Earth System Model, which attempts to represent major components and drivers of the atmosphere, hydrosphere, ecosphere, geosphere and 'sociosphere/anthroposphere' in an integrated manner, is permeating through many scientific and decision-making communities. These models are intended to cast the world we live in, into a holistic quantitative representation of various physical and biological states and fluxes and societal decisions. Traditionally, model development has been conducted in relative 'isolation' or has been 'stove-piped' within individual Geoscience communities. However, the demands placed on many components of the Earth System and the prospects of large scale changes in climate, ecosystems, oceans and landscapes have thrust these communities together as they search for new ways to represent the feedbacks and inter-dependencies between each other. To enable such research in a transformative manner, it is essential to define existing and new model and data structures that facilitate and promote inter-disciplinary collaboration into model development processes. The challenges to creating more coherent modeling systems are multi-faceted though and relate to the rapid pace of scientific and technological development, limited resources for developing and sustaining meaningful collaborations and an existing and enormous diversity in model structures, programming languages, computational platforms and data requirements.

   This EarthCube Roadmap project explores new opportunities for adapting, relaxing and enhancing various model coupling frameworks used in ESMs in order to accelerate broader utilization of these models by the rest of the Earth sciences and decision making communities. In doing so, we advocate here for a 'Geoscience Community Modeling Framework' to accelerate progress on a multitude of aspects related to modeling. This modeling framework must be able to accommodate models of the entire Earth system as well as small parts of it, such as the dynamics of a single valley glacier, the erosion of a particular section of coastline or the hydrology of a small watershed. We believe that this framework can be achieved by building upon and integrating several of the modeling frameworks that already exist while accommodating new frameworks where necessary. It will, however, require cooperation between the various modeling framework projects in order to identify and adopt (or mediate between) common standards for dealing with a variety of interoperability issues including: (1) model metadata, (2) semantic mediation and (3) formats for describing unstructured grids. The core of this Roadmap is a proposed 2-solution strategy that will:

   3. Develop and integrate a set of standards-based protocols for linking component models of a myriad of Earth System processes into a coherent, conservative, multi-scale and high-performance computing enabled modeling system, and
   4. Develop an Earth System Model coupling testbed to facilitate standards-based approaches to ESM component coupling and to support broader engagement and

participation of geoscience domain scientists in community-based ESM development and end-use applications.

In proposing this solution we were guided by several overarching principles which sought to a) optimize engagement and participation across geoscience and relevant end-user communities, b) capitalize on existing and sometimes long-standing model development efforts, c) utilize standards-based approaches to coding and semantic ontologies and d) fully engage current and future computational capabilities. In following these principals we believe this Roadmap is relevant for Geoscience researchers who are involved in the creation or modification of detailed numerical models of earth system processes and want to link their process models to other process models or a larger modeling system. Example processes could be but are not limited to: river flow, atmospheric general circulation, cloud microphysics, groundwater flow, core-mantle interactions, space weather, and ocean biogeochemistry.

Throughout this Roadmap, the term Earth System Model (ESM) will refer to either a standalone model of some process or group of processes or a comprehensive model of the entire global Earth System. We distinguish here between models which are the collections of codes describing myriad Earth System processes and 'modeling frameworks' which refer more specifically to the architectural requirements for linking different process models together. In this sense, 'coupling' can refer to either creating a large comprehensive model from several component models (e.g. atmosphere and ocean to make a climate model) or connecting 2 or more models for some region or sub-system of the Earth.

One common element to nearly all Earth System components is water. Water is ubiquitous, as well as chemically and thermodynamically active within the various spheres of the Earth System. It provides a principle mode of heat and chemical transport, plays an important, mechanical role in the morphology of landscapes and is an essential regulator of global climate. Most importantly, water is essential for life as we know it. As such, representation of its movements and transformations among all components of the Earth System must be executed with exceptional fidelity. However, there remain significant shortcomings in the way in which water is treated in Earth System Models. Presently, only a small portion of the terrestrial Earth science community and even smaller portions of the policy and decision making communities who regular deal with water are engaged in the development and use of fully-coupled ESMs. The reasons for this are manifold but are linked to fundamental incongruities in model conceptualizations and data structures between different scientific and user communities and there are significant technological divides between communities with different communities relying on different standards or conventions for model and dataset development. In recognition of these facts, the EAGER grant under which this Roadmap was funded, was explicitly designed to examine issues related to the coupling of terrestrial hydrology models with other models of the Earth System. At times, throughout the document we draw distinctions between activities and concepts that are relevant to a broad range of Earth System model development efforts and those specific issues that are being addressed more directly through our EAGER grant.

The following sections, respectively in order, describe the Communication, Challenges, Requirements, Status, Solutions, Processes, Timeline, Management and Risk issues associated with our proposed 2-solution strategy. Throughout we attempt to highlight exactly how our proposed solution strategy addresses the needs and challenges we face related to ESM coupling. We also articulate, wherever possible, how we expect community modeling frameworks can lead to greater productivity by making it easy to incorporate existing models and tools from other groups into a new system, and how the ability to plug and play is crucial for repeatability, intercomparison of methods and overall vetting of new scientific and modeling innovations. Finally, we make every attempt to identify the reasons why we feel the proposed solution will optimally expand the community of developers and users of Earth System models.

2. ***Communication:*** *Description of a communications plan with end users, developers, and sponsors, as well as links to and feedback from other EarthCube community groups and concept projects to promote systems integration and accelerate development. Include a discussion of needed interactions with allied fields, agencies, and other related activities (present and desired).*

Overview

The core of the ESM concept group proposal is to provide definitive guidance on how to accelerate model development and coupling efforts. As such, this component of EarthCube has the potential to put powerful new hypothesis –testing capabilities into the hands of geoscientists that should enable new discoveries and pathways to understanding our earth system. To ensure that the new cyberinfrastructure developed is successful in opening ESM to a wider variety of geoscientists, a strategic communications plan must be developed to build ***awareness*** of the project at all stages to **core communities** of both users and developers, to foster ***active engagement,*** early adoption and feedback from these communities in development of ESM infrastructure, and to create a community of ***advocates*** that will assist the project team in disseminating the vision and technologies for adoption by a broad geosciences community. We anticipate that a communications plan will need to evolve, adapt and develop in conjunction with the technical aspects of the project. However, in this Roadmap we lay out a broad ***Communication Strategy*** that can guide the development of these initiatives.

The first step of the strategy will be to **Identify and Engage Core Communities.** We define Core Communities as those who *directly* contribute to or benefit from advances in Earth System Model coupling. In the EAGER phase of this EarthCube project, communication with these core communities will ensure the highest likelihood of producing an informed and useful prototype of optimal modeling coupling strategies for linking a subset of Earth System component models, in this case atmospheric models with terrestrial hydrology models. The ESM team has and will continue to collaborate with and seek input from technical experts and partners and build on past successes of various modeling coupling efforts. The ESM team will also incorporate lessons learned in their prototyping efforts from these experts. After executing its initial pilot demonstration of comparative model coupling strategies, the ESM team will initiate a process of feedback with both users and developers to evaluate the initial experiment and design a path forward. As discussed next, over the EAGER phase of the grant, the communications plan will gradually encompass end users within the Geoscience Community as well as key technical partners.

1. End Users and Contributors within the Geoscience Community

The initial evaluation of modeling coupling strategies and frameworks will be informed by user requirements and functional requirements of the Earth System Modeling community. For the EAGER phase pilot study information and active participation in evaluation will be solicited from researchers who are using atmospheric and hydrology models.

For example, lines of communication will be established through the CUAHSI community (hydrology) and in particular, past participants of CUAHSI's Community Hydrologic Modeling Platform (CHyMP) initiative and current participants in the Integrated Water Resources Science and Services (IWRSS) initiative's National Water Model Effort. The IWRSS initiative is based on a memorandum of understanding between NOAA, USGS, and USACE, and has as one of its goals a "research to operations to research" framework for developing a model for comprehensive water resources forecasting for the agencies. CUAHSI has been supporting academic participation in this effort. While the efforts are distinct, the use cases, functional requirements for comprehensive water resources forecasting, and the experience of the modelers involved in this effort can all inform the ESM coupling concept pilot. Additionally, feedback will be solicited from the operational weather and marine forecasting community, through association with the National Unified Operational Prediction Capability (NUOPC). NUOPC and its research partners are engaged in an active collaboration on scientific and technical issues related to model coupling and associated conventions. Through this connection we will maintain communication and to the extent possible alignment with federal modelers using the Earth System Modeling Framework.

A parallel communication effort will be made to gather input through the National Center for Atmospheric Research's (NCAR) extensive user community and its communication resources. Project team members will solicit feedback from the community of current Community Earth System Model (CESM) and Weather Research and Forecasting (WRF) model users, with special outreach to those involved in past efforts to couple these models to different hydrologic models. From this, we can assemble information on the challenges of the defining coupling requirements, the processes that led to successful coupling, and the science enabled by individual coupling efforts. The community currently most familiar with and using these models have much to contribute on what the challenges and requirements of a more extensible coupling framework would be, as well as a valuable resource for understanding how researchers would use such a framework and what opportunities for new lines of investigation this would enable.

We will also target survey domain scientists who have contributed or used modules of the Community Surface Dynamics Modeling System (CSDMS) project. Identifying contributors is straightforward, and we hope to gather valuable information about usability and adoption of CSDMS as an example of a more general modeling coupling framework. CSDMS interacts closely with its community of 783+ members, 5 working groups and 3 focus research groups and uses this multi-tiered feedback to set priorities and address identified needs. Annual meetings also provide valuable feedback from users and developers.

After completion of the pilot project the results will be disseminated to the broader Geoscience community as well as potential end-users in hydrology and atmospheric science that have not yet been directly involved in model coupling research, but who would benefit from the availability of coupled models if the barriers to use were reduced. The intent of this initial outreach will be to identify lines of extension of the project – including tools that would make the models more available to a wider community and other disciplines and models that could adopt and contribute to the coupling framework.

2. Technical Partners

Model Coupling Frameworks and Projects

As discussed in the 'Status' section below (Section 5), there are a number of modeling frameworks that have received significant attention, development, and funding over the last several years. An important aspect of the coupling exercises proposed in this Roadmap is to evaluate the extent to which geoscientists can work *within* current modeling frameworks, the *commonalities and the differences* between the frameworks, and the pros and cons of each system. Designing an optimal path forward that will benefit the largest number of Earth system scientists will require the technical expertise of those who have been working on these problems for many years. The EAGER phase of this grant will engage the major modeling centers and developers of these frameworks in discussions and collaborations. An important aspect of the communication with these **technical partners**, is that we must develop the *rationale and opportunity* for each of the technical partners.

The *rationale* for developers of current ESM frameworks to participate in this effort is that each of these frameworks is still evolving. There are well-known barriers to adoption of any single framework by a modeler or modeling community - including the level of effort required to make a model framework-compliant and the uncertainty of whether a particular modeling framework will be adopted or supported long term. By evaluating current coupling frameworks for their requirements, current adoption, and current functionality, as well as understanding developers' future plans, this effort will seek to develop a community of practice that provides infrastructure for communication and collaboration.

For example, CSDMS introduced the framework-independent Basic Model Interface (BMI) specifically to remove the requirement that developers adopt or "code to" any particular framework, thereby removing one of the main barriers to developer participation. In theory, a BMI-enabled model should be easy to wrap with whatever component interface a particular modeling framework uses. In practice, however, the BMI needs to be evaluated for its completeness, its overhead, and its applicability to a range of model components. The same sort of evaluation would need to occur for other model coupling frameworks as well.

The *opportunity* provided by this collaboration is that frank discussion of technical challenges, approaches, successes and failures, will not only help this project in its aim of producing guidance for ESM coupling approach(es) that will most benefit the geoscience community, but also will help advance individual modeling frameworks by identifying factors such as (1) key functionalities of individual frameworks that indicate (2) specific applications, communities or use cases for which a particular framework is optimal, (3) strategies or design decisions made in one framework that may be ported and/or appropriate for others, and (4) the extent to which requirements for individual frameworks can be standardized, which would ultimately lower the barriers for model compatibility for all frameworks.

Data Providers

For the proposed coupling architecture to ultimately benefit the largest swath of Earth System scientists possible, the architecture developed must eventually include integration into heterogeneous workflows and cyberinfrastructure support. Particularly, if the architecture is to expand the *users* of earth system models beyond the current community (in addition to expanding the ability of the current model community to interact with each other, each others' models, and different frameworks), then the communications plans must expand out beyond the model development community and recruit partners that provide other aspects of a complete, end-to-end system to enhance hydrologic and atmospheric modeling. One of the most important aspects of this is data. Thus, this project will engage data providers and developers of data software solutions to best understand how work done in this project can integrate with these other aspects of end to end modeling. These partners include (1) OPenDAP, which develops open source software for serving diverse types of data, (2) Unidata, which provides data services, tools and cyberinfrastructure for the atmospheric science research community and (3) CUAHSI HIS, which provides data services for the hydrologic community.

Other EarthCube Groups

Collaborations have already began between several of the EarthCube Community Groups and Concept Groups, as it is recognized that this project is a piece of a larger system of cyberinfrastructure that will serve the geosciences and transform the way earth system science can be done. The design of EarthCube facilitates interaction amongst the group, and this project has maintained open lines of communication in the following ways:

- Identifying Liaisons with other concept grants and community groups
- Core team members of this group participating in workshops and webinars of other EarthCube groups
- Inviting members of other EarthCube groups to present related work

As the community groups evolve and as concept groups move forward on their projects, we expect to identify and pursue specific avenues of collaboration and intersection by working closely with other groups. Examples of instances where we have identified specific overlapping interests or areas for collaboration include:

- **Workflow Community Group:** End users of a coupled Earth System model will use the model within their own explicit or implicit workflow. Understanding how Earth scientists and potential model users design their workflow process could impact the functional requirements and design decisions of a community modeling architecture. Additionally, understanding different formal workflow systems and their requirements will be important in understanding how the results of this concept grant will impact end to end modeling.
- **Data Discovery, Mining and Access (DDMA) Community Group:** As stated above, the availability of data, in the correct forms, and properly documented, greatly increases the utility and usability of models, regardless of the framework. Understanding current data access and

delivery services and standards will be an important aspect to the eventual design of a community coupled earth system model architecture.

- **Brokering Group**: Some existing modeling frameworks, such as ESMF, have the ability to represent modeling components as interactive web services. Hence, various models or modeling components have the potential to be incorporated into workflows and model coupling architectures that include brokering services and potentially other data services.
- **Layered Architecture Concept Group:** It may be feasible to connect an integrated modeling system such as CSDMS, or those explicitly mentioned in this Roadmap, with an integrated data system such as iRODS as parts of a layered architecture for EarthCube. While iRODS provides robust and scalable data management capabilities and support for simplified workflows, it does not currently support component-based modeling (i.e. it has "logical namespaces" for several types of entities, but not yet models) or provide the semantic mediation needed for model-to-model or model-to-data interoperability. As such, this ESM coupling Roadmap could simultaneously contribute to and benefit from progress in the development of layered architectures.
- **Cross Domain Interoperability Concept Group:** The issues the 'inter-operability' group are addressing (e.g. describing and preparing data and models across domains for reuse by different communities) are directly relevant to core aspects of model coupling - including how to sufficiently describe model components and evaluate their suitability and/or interchangeability in different situations.

Federal Agencies and Third Party Partners

Understanding, predicting and managing resources in light of earth system processes are major foci for a number of federal entities, including the Department of Energy, the Department of Commerce (e.g NOAA), the Department of Defense (e.g. the US Army Corps of Engineers), the Department of the Interior (e.g. USGS), and the Environmental Protection Agency. Many of the current modeling systems have primary funding provided by different federal agencies, and development of new modeling systems and technologies are of interest across many of these entities. This project will communicate with federal agencies and other third party partners that are employing or interested in coupled earth systems models for their operational needs. The communication strategy will focus on highlighting the *rationale* for these groups to be involved or engaged with this project and the *opportunity* to design mutually beneficial interactions.

Many federal agencies have invested significant resources in modeling frameworks, and/or have missions that may be best served by adopting a coupled Earth-system viewpoint. Additionally, federal agencies have a positive obligation to share data, model results, and promote interoperability amongst agency systems. The academic community represents a valuable knowledge resource in meeting these goals. Agency scientists and administrators often recognize this and seek out academic partners or contracts to meet specific project goals. A large scale effort on the order of Earth Cube should be of interest to a wide variety of federal and other partners because of the critical mass it provides and the opportunity to collaborate and harness the best science and personnel in ways that are mutually beneficial.

The project team has already reached out to several potential partners for advancing our understanding and the potential applications of better understanding the coupled earth system. Several of the prominent modeling frameworks discussed above are funded by federal agencies and led by agency scientists. The project team will work to continue the dialogue began at our group's pre-Charrette workshop at NCAR in May of 2012, and strive to develop recommendations and best practices that span across the various frameworks and functional requirements of different systems.

In addition to working with technical partners from these efforts, we will work to engage agency administrators and specific projects and initiatives that would benefit from the *applications* of coupled model frameworks (in addition to those *developing* them). One example of this is the inter-agency initiative **IWRSS** (Integrated Water Resources Science and Services), which is undergoing an exercise to scope, design and implement a national model for comprehensive water resource forecasting to meet the needs of the involved agencies (NOAA, US Army Corps of Engineers, and the USGS). CUAHSI is partnering with the agencies to facilitate academic involvement in this exercise, and members of this ESM Concept Group team have been invited to participate in the initial scoping exercises. While the IWRSS water model will focus on meeting operational and agency mission needs, to meet these needs requires the technical and science advances for which this concept grant is aiming. The vision of the IWRSS water model is a "research to operations to research" framework where academic and agency scientists contribute to an operational model in a standard research to operations pathway, but that the system is designed in a way that the model structure, data and results are "open" to researchers to further test and recommend innovations to the model and framework through a formal system. This vision overlaps with many of the themes running through Earth Cube, and it is our belief that the "comprehensive" model could benefit from the results of the exercises we are proposing here, including lowering the barriers for individual modelers to participate in different frameworks and for scientists wishing to use various models for scientific investigation. With CUAHSI involved in both the IWRSS initiative and as a formal partner on this grant, the lines of communication have been readily established. Additional lines of engagement between other similar partnerships or through other avenues will need to be forged as well.

Post-EAGER Communication Strategy

The EAGER phase of the communication strategy focuses on identifying and engaging potential technical partners and inventorying the needs, requirements and characteristics of the potential end user communities, and designing a careful integration strategy with other Earth Cube groups. These interactions are meant to ensure that the EAGER phase of the concept grant produce concrete results and recommendations that can inform future developments and demonstrate the utility to the technical partners of their involvement.

Post-EAGER, we recognize that a broader communications strategy that builds upon an increasing technical effort over time will be needed. While the communications strategy will evolve with the project, we recommend several key aspects of a communication infrastructure be established over the three to five year horizon as the concept expands to a full earth system model coupling infrastructure.

1. Continued Community Resources maintained by Earth Cube or another Entity

In the initial phase of the grant, the ability to share discussions amongst a community site or email list has proven to yield valuable contributions from those beyond the initial core team. Additionally, the ability to see cross-connections between this and other Earth Cube Concept grants has strengthened the proposed plan. Most importantly, the ability to receive formative and summative feedback from potential end users during the different phases of the project provide a significant advantage over the way modeling framework development usually occurs. This interactive framework will help develop the most appropriate use cases, pilot exercises, and ensure that developments serve the needs of the end users and model developers *in the geosciences*, and will be an important aspect of ensuring long term adoption by the earth science community.

2. Community Workshops to Present Results and Gather Additional Feedback
After the completion of the EAGER phase, we anticipate producing a set of recommendations as well as the results of the pilot project that will inform the development of the next phase of model coupling infrastructure. Given the short time frame of preparing for the June 2012 meeting, the initial workshop held May 22, 2012 was deemed to be exceedingly productive in establishing connections between the current modeling frameworks and engaging a number of atmospheric and hydrologic modelers in beginning to evaluate what more is needed. These small workshops should be continued as the project evolved. Semi-annual workshops that focus on reporting recent changes and status of current modeling frameworks to the community, and then shift toward providing target recommendations from the user and developer community for the next phase of the project, should be established to provide targeted benchmarks for development of the model coupling framework and iterative feedback to improve the concept.

3. Establishment of a Community of Practice for model developers
Recommendations and findings of this project will only be as useful as the extent of their adoption. Instilling a community of practice among model developers has been identified as a goal in other community modeling initiatives (e.g. CUAHSI's Community Hydrological Model Program (CHyMP) initiative and would create a central entity for providing information and training on model coupling standards and recommendations, a central location for model developers to share information about their models and for end users to find information on different models and frameworks to help them evaluate what best meets their science needs and can be used to establish a governance framework for drafting, reviewing and recommending model coupling standards for the community. The Community of Practice would also interact with other aspects of EarthCube, and advocate for developing standards or best practices for model metadata descriptions, to ensure that models that are compatible with various frameworks are documented in sufficient detail to be used and coupled consistently and correctly by a wide variety of Earth scientists.

4. Formal Liaisons with Data Providers
In the EAGER phase of the project, the most direct connections have been made thus far with the technical experts developing state of the art model coupling strategies. However, critical links must also be made with data providers and centers to ensure that the infrastructure being developed meets the needs of users and developers to more easily and effectively execute model simulations to meets their science goals. A critical piece of a simulation exercise is the data that drives the simulations. We know from past modeling exercises and initiatives, that access to high quality, complete sets of the full spectrum of data needed to run coupled earth system models, in a form and structure that can be incorporated into a particular model or class of models, is often cited as a critical limitation of our current capabilities. Web services, integrated workflows, data discovery and access technologies, and other cyberinfrastructure advancements being proposed or examined by others in the Earth Cube

community may help address some of these limitations, and several of the leading model framework developers have worked directly with data providers and technologies to explore interoperability of specific frameworks with specific data delivery technologies. As this project advances, there is an opportunity for the Community of Practice discussed above to serve as the focal point for the earth systems model community to aggregate experiences, articulate current limitations and needs, and advocate for specific enhancements to data availability and technology needed to enhance our ability to model the earth system. In the second phase of a model coupling project, we would anticipate greater infrastructure and human capital that could be dedicated to formal outreach and coordination with different earth science data providers and centers (many of whom are heavily involved in Earth Cube) to promote more seamless integration of data into model simulations.

3. ***Challenges:*** *Description of major drivers, trends, and shifts that are or could impact the focus of your concept group, including but not limited to changing technology, adoption culture, and community engagement.*

The Earth System modeling community faces numerous challenges, both social and technical. Climate variability and change and the ever-increasing human population are key drivers that have led to increased water and other resource stresses across many sectors and geopolitical lines. In the future, these factors will likely continue potentially affecting income disparity and societal capacity divergence and inequality. As society attempts to address these issues, there is an increasing demand for accurate, easy-to-use predictive models. Paradoxically, the demand for understanding is not always commensurate with the demand for ease of use. However, in order for the results from models of the Earth System to be accepted broadly and to lead to sound societal policy changes, these models also need to be open, transparent, validated and easy to compare. There are several technical challenges which must be acknowledged and addressed in order to do so. Important trends such as; (1) increasing specialization within the earth sciences, (2) models that work on one piece of the earth system without due consideration to interdependencies on other Earth System components, (3) proliferation of technology through user and value chains and (4) Moore's law of exponentially increasing computational capacity all must be fully acknowledged and appreciated. Each of these factors make it difficult to coherently assess and design modeling frameworks that need to be both robust to new scientific and computational innovations while remaining accessible to the broadest range of code developers and model users.

Institutional challenges also play a big role in determining the pace of model advancement. Institutional lethargies can foster resistance to change or advances from outside parties (i.e. the 'not-invented-here' culture) and many coupling solutions have been developed in small groups which may lack a broader perspective needed to develop more comprehensive solutions. Furthermore each institution generally needs to follow a path that is most closely defined by its institutional mission. There is also a difficult balancing that must occur with many models between the need for progressive advancement of modeling tools and the need for stable, regulatory/management tools. These institutional barriers have led to markedly different coding structures, differences in tolerances for errors and uncertainty, and shortcomings with respect to education, training and re-training of model development and user communities.

Perhaps the biggest technical issue we face as a modeling community is the issue of "semantic mediation" via a set of "standard names" that is supported by underlying metadata (like units, assumptions, how quantities are measured, etc.). We only have a few domain-specific controlled vocabularies at this point (e.g. CF Standard Names and CUAHSI-HIS Controlled Vocabulary) and they were not designed to serve as a "lingua franca" for sharing state and flux variables between different models and databases. This is a cross-cutting issue that requires community participation and buy-in. It is part of the whole "semantic web" issue/goal that must be solved in order to create seamless, automated systems and workflows that work as intended while requiring minimal effort from users. The role and importance of metadata --- for models, databases, workflows, etc. --- is very closely related and there are a variety of metadata standards that need to be considered including XML, SQL, RDF and various OGC standards.

A host of other challenges include: (1) uncertainty tracking and analysis in modeling workflows (for which major advances are still needed), (2) the need to calibrate new models that have been created by assembling components, and general tools for model calibration, (3) a better and more general understanding of numerical instabilities that sometimes arise when coupling models that are individually stable, (4) automated use of model metadata to alert users to model component coupling opportunities as well as potential mismatches/conflicts in terms of their individual assumptions, applicable scales, etc., (5) a seldom-acknowledged tension between the typical timescales for a) funding of cyberinfrastructure projects (3 to 5 years), b) maturation and adoption of new cyberinfrastructure (4 to 10 years) and c) the need to maintain cyberinfrastructure once it has matured and been adopted (4 to 20 years).

Given this rather daunting summary of challenges it is only realistic to consider a solution for Earth System Model coupling which is adaptable and evolutionary in the sense that the solution can change over time in response to the demands of the community relying on it rather than presuming to design a singular end-to-end solution that will last in perpetuity. This final point is one of the key motivating factors behind our proposed 2-solution strategy.

4. ***Requirements:*** *Process(es) to be used to get the necessary technical, conceptual, and/or community (i.e., end-user) requirements at the outset and during the life of the activity, including approaches to achieving community/end-user consensus.*

In this section we articulate a set of requirements or ideal attributes for developing new modeling frameworks which will facilitate efficiencies in Earth System component model coupling. These requirements were synthesized from the experiences of the ESM concept team and through direct engagement of many ESM developers through the Roadmap drafting process. As such they represent a broad range of perspectives and needs related to increasing access to and usability of ESMs. The communication processes by which these requirements are vetted by an even broader community are described at length in Section 2, above. The section is broadly divided into two parts, one section summarizing the community modeling philosophy which it is felt future ESM development efforts should follow and another section listing a set of 'functional attributes' of ESMs and their coupling architectures.

4a. Adherence to a 'Community Modeling' Philosophy

One principal goal of this modeling Roadmap is to advocate for the concept of 'community-based' modeling. In this section we provide a working definition or philosophy of what constitutes a community model. It is not intended that this definition serve all purposes for all model development efforts. While some of the criteria posed below could be applicable to broad categories of models, many others may be more nuanced or specific to the domain of Earth System process models as that is the topic of this Roadmap. As such, this definition posed as an attempt to constrain the application of this term of art as it is relevant to and used within this specific document. Therefore, we define the following 4 attributes as ideals of a 'community model':

1. The source code which describes the Earth System processes conceptualized within a particular model should be open source and freely available. Additionally, to the greatest practical degree possible, functional libraries, compilers and execution scripts should all be freely available and manipulatable by the broader community of developers and users.
2. The model development process including the formulation of modeling experiments and model assessment projects needs to be open to a broad community of developers and knowledgeable stakeholders. This participatory process must allow for presentation, debate and analysis of new modeling concepts and formulations. However, this open development framework needs to be arbitrated by some supported authority or team of 'core personnel' to ensure long-term viability and structured evolution of the modeling system. Additional detail on what that arbitration process may look like is discussed in Section 9, Management.
3. The modeling system should be supported by some entity such that thorough and current documentation is available. Training, in the form of live and/or online model tutorials should also be available, though, particularly in the case of live training sessions, such training may require participant costs. An online user support network, either via mailing lists, support desks, chat rooms or other means is available for developers and users having well-defined, technical questions.
4. Standardized and community-defined modeling use cases and assessment benchmarks in terms of both computational performance evaluation methodologies and accuracy of model solutions should be developed and openly available. These modeling benchmarks provide clear and quantitative, albeit somewhat subjective, criteria for demonstrating and documenting the performance of new models or enhancement to existing models. Such model benchmark problems should incorporate state-of-the-science understanding in the definition of the modeling problem as well as in the verification and interpretation processes.

While it is admittedly difficult in practice to achieve all three of these ideals they are proposed here as ideals to be aimed for in the context of engaging the maximum number and diversity of scientists and practitioners in the development and use of Earth System models.

4b. Earth System Model Protocols with Explicit Emphasis on Extensibilty

Modern computational models of the Earth System need to possess several important functional attributes so that they may successfully serve their designed purposes. These desired attributes range from basic scientific criteria to criteria that are either computationally oriented or are relevant for the range of societal needs for which the models were developed. For the context of this Roadmap we list these desired attributes as follows:

- **Conservative formulations:** Model codes must be conservative with respect to mass and energy transfers within and between model components. Mass and energy conservation provide fundamental constraints on our conceptualization of nearly all earth system processes. While certain applications of various ESM components may not require complete conservation, we advocate here that to better understand the interactions between and dependencies of various Earth System components between one another closure of all energy and mass balances must be maintained to a fine degree of model precision.
- **Extensibility:** The coupling architectures within any respective ESM must be readily extensible by a knowledgeable but diverse development and user base. Understanding of Earth System processes is increasingly complex and requires the interactions of scientists and practitioners from a vast array of disciplines or 'domains'. If a particular model coupling architecture is overly limited in the types of model components that can be joined together, or if the process of adding new components or maintaining model version control is too onerous development of the system stalls and innovations in process understanding diminish. Therefore, to foster an optimal balance between model fidelity and model applicability we advocate that coupling architectures for Earth System models be robust to a wide variety of component additions and subtractions and be very clear in their programming protocols.
- **Multi-physics capabilities:** Conceptualizations of many Earth system processes are presently evolving. As such, many different and equally-plausible mathematical formulations of these processes have been created by disciplinary/domain scientists. For instance, numerous formulations exist for the process of flow through porous media such as occurs with variably saturated sub-surface flow. While fundamental equations, such as 'Richard's Equation' can be used, there are many instances where alternate formulations must be employed (such as flow through karstic aquifers). To accommodate such process uncertainty we advocate that Earth Systems model coupling frameworks be able to robustly enable multiple-physics parameterizations for similar processes (i.e., 'multi-physics' capabilities) in addition to a single suite of physics. While it is important to maintain the fidelity of many single-suite physics models, we suggest that employing multi-physics capabilities will enable more robust hypothesis testing across scientific domains and also provide quantitative means to account for model conceptualization uncertainties.
- **Multi-scale functionality:** Similar to the previous point of enabling multi-physics capabilities in Earth System models, we also advocate for ESMs to enable 'multi-scale' modeling frameworks. We define 'multi-scale' as the ability to simulate different Earth System processes on different spatial or temporal scales and to efficiently and conservatively map model states and fluxes

between model grids of different spatial resolutions. Many recent development efforts in ESMs such as those occurring within the NCAR CESM are already developing multi-scale modeling capabilities between and within various model components, such as atmosphere, land surface and ocean components. While performing conservative spatial regridding, multi-scale modeling tools must often mediate time-stepping differences between modeling components. We support this development effort and encourage new coupling frameworks to be developed that further enhances these capabilities.

- **High Performance Computing:** Nearly all current ESMs utilize high performance computing architectures which permit the simultaneous execution of model calculations on multiple processing cores. As process diversity in model formulations increases, spatial resolution decreases and model simulation durations increase, the computational demand required to execute ESM code increases, often exponentially. While many components of Earth System processes were not developed within parallel computing architectures, we advocate for their implementations within ESMs to be performed in such a manner as to permit full-parallelization. This activity often requires skillsets beyond the typical domain scientist capabilities and usually requires significant partnership with software engineering, scientific computing and other cyberinfrastructure domain areas of expertise. Nevertheless, to realize the overall vision of an ESM with maximum utility to domain scientists and societal applications alike we encourage these collaborations and the associated development of parallelizable code.

- Cross-Platform Portability: To the greatest degree possible, community model code should be ported or at least be portable to a wide variety of computational architectures. However, this isn't to suggest that Earth System Models be unnecessarily simplified either in model conceptualization or cast in reduced spatial and/or temporal resolution so that they can run on all platforms. This portability attribute is intended to suggest that model developers think broadly about the various classes of computational systems available to run a suite of Earth System model components and design code that is reasonably adaptable to such architectures. In the context of current Earth System Models which possess a significant atmospheric component (i.e. weather and climate models) this necessitates the use of parallel processing architectures at a minimum and, likely so-called High Performance Computing systems which are capable of executing model code on thousands of processing cores simultaneously. Hence, to achieve this attribute, new Earth System component model development efforts should incorporate parallel coding capabilities as early as possible into the code design and development process.

- **Language Interoperability Support:** New integrated ESMs should have significantly expanded cross-language inter-operability. Geoscientists and end-users alike often have limited access to major computing facilities and architectures on which current ESMs are developed and operated. While it is unlikely that many core components of ESMs will be required to run on non-HPC architectures, many of the ancillary models that interact with core physics components do. As such new coupling architectures should seek to broaden access to ESMs by expanding porting efforts to new computational platforms and should incorporate the use of multiprogramming language interoperability applications. Different programming languages offer different advantages and with the advent of a powerful, language interoperability tool like Babel, there is no reason to restrict programmers to using only a few of the available languages. While code written in languages like C and Fortran often runs faster, code written in a modern, object-oriented language like Python, C++ and Java offers myriad benefits in terms of portability, maintenance, elegance, readability, development speed, code organization, etc. Babel provides the "glue code" that allows a component written in a supported language (C, C++, Fortran [all years], Java, Python) to efficiently retrieve data from a component written in any of the other supported languages. Babel uses an elegant hub-and-spoke design based on SIDL (Scientific Interface Definition Language) and could be extended to support additional languages such as C# (used for many commercial applications on Windows PCs) and MatLab (used by many scientists for smaller modeling projects). A key benefit of using a language interoperability tool is that it creates a more inclusive and flexible environment; more developers are able to contribute a wider array of models and tools. Both CSDMS and MCT, described below in the next section, utilize Babel. We expect that Babel could play a very important role with regard to the model interoperability aspects of EarthCube. Incorporation of these criteria would serve to significantly broaden access of ESMs to a wider range of geoscience model developers as well as long-tail ESM end-users.

- Semantic Mediation and Units for variable names: As diversity in model processes, structures, and architectures expands there are challenges associated with the use of different variable names and units for similar terms passed between models. To the greatest degree possible, standards-based data models and data flow structures should be utilized. This attribute implies that using standards-based data models is an important aspect of the model development process. Incorporating such considerations from the model design phase all the way through to analysis, visualization and archiving stages of the modeling workflow can offer tremendous efficiencies in terms of defining and tracking data attributes. It also allows for the maximized use of pre-existing analysis and visualization software tools in the model pre-processing and post-processing phases. Furthermore, use of standardized data models also facilitates improved data discovery and mining capabilities both leading up to and following a specific modeling experiment. However, newly developed component models and model coupling tools will often need to mediate differences between component model variable names and units. When it arises, this 'semantic mediation' process needs to capitalize on evolving standards in data naming and units conventions such as those utilized in the 'netcdf' convention, for example. Promotion and use of such data standards within component model code and more generic coupling architectures will minimize secondary, ad-hoc semantic mediation efforts.

As an outgrowth of the EarthCube community building effort, we also acknowledge and advocate an additional requirement related to strengthening the inter-connectivity between the traditional geoscience modeling communities and sister cyber-infrastructure communities. As discussed in the

Communication section above (Section 2), several other cyberinfrastructure groups are addressing key issues in the standardization of workflows and in the development and mediation of data architectures that are critical the overall Earth System modeling process. To this effect the following architecture framework has been proposed as a template for guiding connectivity between groups addressing various cyberinfrastructure components. The architecture framework (below) that has been proposed is dominated by web services which seek to interconnect developers and users of geoscience information. While many components of the proposed architecture are directly relevant to Earth Systems modeling, there are some deficiencies which are specific to current modeling activities. Specifically, because of the very significant computational and memory demands of current Earth Systems models, it is imperative that proposed workflows and architectures account for the need to accumulate relevant information streams onto a single High Performance Computing (HPC) platform. In the present diagram this need is visualized as the 'brown cloud' backdrop. While possible, it is presently impractical to execute large, multi-scale, integrating model runs across a web service platform, mainly due to the massive amount of data transfer that must occur during runtime but also due to subtle differences in machine precision and other system specific features. Therefore, while this Earth Systems Modeling Roadmap endorses the proposed architecture framework for connectivity between the various components of the EarthCube enterprise, we note that there are a set of issues specific to Earth Systems models and their associated computational requirements which are not fully explicit and, as articulated above in this requirements section, will need to be addressed in ongoing and future community discussions.

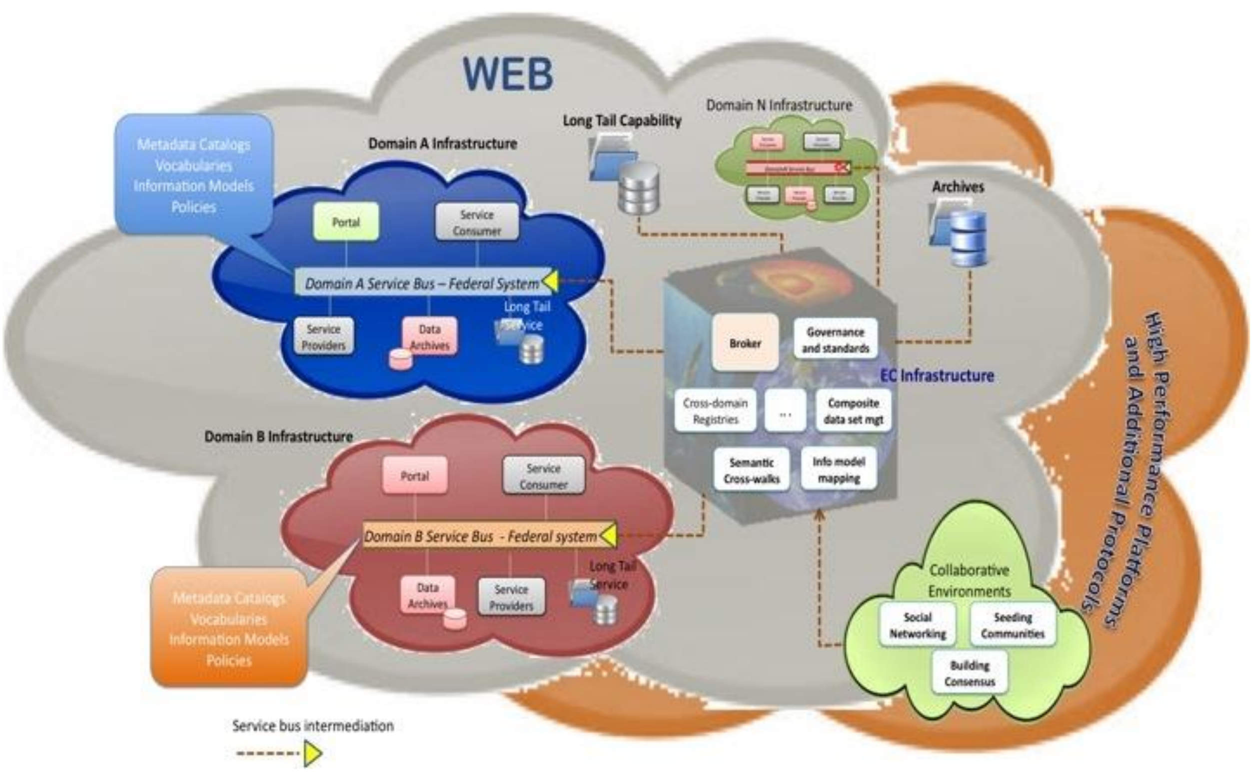

Proposed EarthCube Architecture Framework

5. ***Status:*** *Description of the state-of-the-art within the topical area of your roadmap. This should include approaches and technologies from geoscience, cyberinfrastructure, and other fields, the public or commercial sector, etc. that have the potential to benefit the EarthCube enterprise.*

In this section we provide a technical summary of the major model coupling frameworks that currently exist. We have attempted to standardized the entry for each model so that they can be readily compared in a qualitative sense. Of course, in forcing a descriptive structure for all frameworks we acknowledge that we may not be able to most completely detail all model framework features and uses to their fullest extent.

Current Modeling Framework Projects
The modeling framework projects that have attracted the most interest over the past 5 to 10 years are: (NOTE: This listing is not in any order of priority and suppositions to that effect must be avoided.)

(1) CCA (Common Component Architecture)
- community/domain: scientific, high-performance model users/developers
- funding source: DOE
- platform: mostly HPC and Unix-based (including Linux, Macs)
- uses: GNU compilers, MPI, Boost, SIDL, Python, etc.
- provides: framework standards, tools: Babel, Bocca, Ccaffeine
- GUI: CCAFE-GUI, and can now use CMT from CSDMS.
- URL: sourceforge.net/projects/cca-forum, www.cca-forum.org

(2) CSDMS (Community Surface Dynamics Modeling System)
- community/domain: Earth surface process model users and developers
- funding source: NSF, industry partners, other
- platform: HPC Linux cluster (currently), other Unix, Mac (soon)
- uses: CCA toolchain (Babel, Bocca, Ccaffeine framework)
  ESMF/ESMP regridder, NetCDF,
  VisIT multi-processor visualization software (DOE),
  GNU compilers, Python/NumPy, Java, svn, etc.
- provides: Tools for converting open-source models written in different languages (C, C++, Fortran, Java or Python) by different authors to plug-and-play components with tabbed-dialog GUIs and help pages.
  A repository of 205+ open-source models and tools.
  A collection of 55+ plug-and-play model components.
  A modeling framework and GUI for coupling model components to create new models.
  BMI (Basic Model Interface) as a developer target and CMI (Component Model Interface) for use within the framework.
  Service components for unit conversion, netCDF output, spatial regridding (via ESMF regridder or OpenMI regridder), etc.
- GUI: CMT (CSDMS Modeling Tool), light-weight Java client,

cross-platform. Provides graphical model assembly, configuration, saving, execution, job management, etc. Each model component gets its own GUI dialog and HTML help.
URL: csdms.colorado.edu

(3) ESMF (Earth System Modeling Framework)
- community/domain: domains used in include space weather, weather, climate, coastal, for research and operational codes, regridding tools used by data community (NCL, UV-CDAT, etc.) and other frameworks (OASIS, CSDMS), about 4200 downloads.
- funding source: DoD, NASA, NOAA, NSF
- platform: tested routinely on 36+ platform/compiler combos, HPC, Mac, Linux
- uses: NetCDF, C++, Xerces (XML), minimal external dependencies
- provides: 75+ components in Fortran or C, 12+ coupled modeling systems, confirmed scaling out to 16K+ processors, optimized mesh, logically rectangular, or observational data structures, separate parallel regridding tools, Python interface option, switch for web service interfaces on components, automated generation of metadata in Common Information Model (CIM) XML schema
- GUI: no native GUI; integrated into GUI environments such as the Purdue CESM portal; integrated into workflow GUIs such as Kepler using the web services option
- URL: www.earthsystemmodeling.org

(4) CESM (Community Earth System Model)
- community/domain: mostly global climate model users/developers
- funding source: NOAA
- platform: mostly HPC and Unix-based (including Linux, Macs??)
- uses: ESMF, MCT, lots of Fortran
- provides: a global climate model, mostly ocean-atmosphere but with increasing levels of land surface processes and hydrology
- origin: CCSM (Community Climate System Model)
- GUI: CCSM-GUI (PerlTK)
- URL: www.cesm.ucar.edu

(5) MCT (Model Coupling Toolkit)
- community/domain: scientific, high-performance model users/developers
- funding source: DOE (Climate & Environ. Sci. Div., Office of Bio. & Environ. Res.)
- platform: many HPC and Unix-based (including Linux, Macs)
- uses: Fortran90, Babel, (not MPI)
- provides: Low-level coupling solutions for HPC models
  C++ and Python bindings (using Babel)
- GUI: not applicable
- URL: www.mcs.anl.gov/research/projects/mct/

(6) OpenMI (Open Model Interface)
- community/domain: Windows PC model users/developers

funding source: European Commission, commercial software partners: Deltares, HR Wallingford, DHI (Danish Hydraulics Institute)
platform: Window PC, some Java support
uses: Microsoft Windows .Net framework, C#, IronPython ??
provides: an object-oriented, model coupling interface standard, SDKs for .Net and some Java
GUI: Configuration Editor (included in Windows SDK)
URL: www.openmi.org

(7) OMS (Object Modeling System)

community/domain: agriculture and environmental models
funding source: USDA / ARS (Agricultural Research Service)
platform: PCs with Java (Windows, Mac, Linux)
uses: JDK 1.6, GNU compilers, multithreading
provides: About 114+ model components.
Ability to wrap Fortran models (also C and C++ ?).
Model assembly, execution and visualization.
GUI: Java-based GUI (Netbeans ??)
URL: javaforge.com/project/oms
origin: MMS (Modular Modeling System), funded by USGS

(8) FRAMES (Framework for Risk Analysis of Multi-media Environmental Systems)

community/domain: primarily used for environmental models (fate and transport, risk assessment), although designed to allow disparate models to link and communicate
funding source: EPA-ORD, EPA-ORIA, DoE-PNNL, NRC, DoD/ERDC
platform: PCs (x86 Windows)
uses: Visual Basic 6.0, Visual C++ 6.0, JDK1.7. Single threaded.
Supports models written in FORTRAN, Java, .Net, Python
provides: * several modeling domains, each containing between 5 and 17 linked models (130+ model components in total)
* Ability to wrap .NET, FORTRAN, C, C++, Java, Python models (other languages/compilers can be added quite easily)
* SuperMUSE software containing extensive monte carlo based uncertainty/sensitivity analysis functionality linkable to any model(s) & software for managing embarrassing parallel modeling system runs on PC clusters
GUI: Visual Basic (series of editors to facilitate model registration, units conversion factors, model I/O dictionaries, model linking)
URLs:
www.epa.gov/athens/research/modeling/3mra.html<http://www.epa.gov/athens/research/modeling/3mra.html>
http://mepas.pnnl.gov/EARTH/index.stm<http://mepas.pnnl.gov/EARTH/index.stm>
http://mepas.pnnl.gov/FramesV1/index.stm<http://mepas.pnnl.gov/FramesV1/index.stm>
http://mepas.pnnl.gov/FramesV2/index.stm<http://mepas.pnnl.gov/FramesV2/index.stm>

* https://iemhub.org/resources/133*<http://mepas.pnnl.gov/FramesV2/index.stm>

(9) BFG (Bespoke Framework Generator)

community/domain:	climate and climate impacts
funding source:	E.U. Seventh Framework Programme - IS-ENES, IS-ENES2
platform:	Generates code that runs on HPC platforms
uses:	Python internals (previously XSLT)
provides:	Generative programming tool which takes in metadata describing a coupled model and produces wrapper code using a particular software infrastructure (OASIS, TDT, web service interface, new ESMF)
GUI:	not applicable
URL:	http://www.cs.man.ac.uk/cnc/projects/bfg

In compiling this list it is natural to attempt to draw out basic similarities and differences between architectures. Such an activity is in fact a valuable learning tool for developing existential or vicarious experience about what seems to be working and what are common challenges. Here we summarize some of the major similarities and differences.

Similarities Between Modeling Framework Projects

- most model coupling interfaces use an Initialize, Update, Finalize (IUF) pattern
- all have needs for semantic mediation
- most need to accommodate different time-stepping schemes
- increasingly most need to accommodate different computational grids
- most need to accommodate different data types
- all have large user/developer communities

Differences Between Modeling Framework Projects

- supported computing platforms (e.g. HPC vs. PC, OS, etc.)
- types of data structures supported (provides interpolation weights, calendars, metadata services, etc.)
- level of testing
- level of user support
- ease of use and effort required for initial adoption
- structure and maturity of governance
- level and stability of funding
- scalability and performance
- supported programming languages (language interoperability)
- availability of web service interfaces
- availability of a graphical user interface (GUI) and its capabilities
- support for parallel decomposition of grids

An Example of a Framework-Independent Standard to Promote Interoperability

Here we provide one example which illustrates what the present status of a framework-independent architecture can look like. It is offered for illustrations purposes only.

One of the key achievements of the CSDMS project has been the development of an innovative, two-level wrapping process (BMI/CMI) that greatly simplifies the process of converting contributed models into interoperable, plug-and-play components. Model contributors are asked to make relatively small changes and additions (e.g. functions that describe their model's attributes in a standard way) to their source code to provide a Basic Model Interface or BMI. BMI implementation is noninvasive and straightforward --- it requires no calls to CSDMS code and no knowledge of CSDMS framework concepts or protocols. By design, BMI provides all of the model information (grid type, information on input and output variables, etc.) that is needed by a second-level wrapper that converts the model to a CSDMS component. The second-level wrapper provides a Component Model Interface or CMI that enables coupling to other CSDMS components and automatically calls service components when needed to accommodate numerous differences between models such as programming language, computational grid, time-stepping scheme, variable names and units. Service components provide additional added value such as output to NetCDF files, unit conversion and spatial regridding. By design, BMI allows the same CMI wrapper to be used for every model written in a given language. This greatly simplifies and reduces maintenance associated with the wrapping process and reduces the burden on code contributors. A key benefit of the BMI/CMI approach is that it provides a possible mechanism for sharing models between frameworks. There is nothing framework-related in a BMI-enabled model and yet the BMI interface allows a caller to retrieve anything it needs for deployment in a framework. We expect it will therefore be straightforward to wrap a BMI-enabled model to provide an interface other than CMI, as would be needed for use in another framework like ESMF, OpenMI or OMS.

6. ***Solutions:*** *Process for the identification and comparison (pros and cons) of approaches and technology solutions that will contribute to the EarthCube goal of satisfying current and future research needs of the geoscience end-user. In addition, the process and criteria used in the evaluation should be described.*

In this section we specify what high-level courses of action are required to accelerate progress on the development of and accessibility to Earth System models by the broader Geoscience and end user communities. Statement of our two over-arching solutions are and justification for how these solutions were arrived at are elaborated on here. The specific processes by which these solutions will be evaluated is elaborated upon in the following section, Section 7.

Earth System Models today are required to serve an ever-widening number of roles in scientific research, operational forecasting, resource management and decision making. Though they have emerged from climate models the breadth of processes engaged now clearly transcends much of the geoscience domain. As such, any evaluation of ESM capabilities or assessment of their structure and function from a cyberinfrastructure perspective must incorporate relevant and forward looking applications of the modeling systems. The first solution to the Earth System Model coupling challenge as articulated above that we propose is to develop and integrate a set of standards-based protocols for linking component models of myriad Earth System processes into a coherent, conservative, multi-scale and high-performance computing enabled modeling system. We envision this solution evolving, not as a single monolithic system, but, instead, a layered architecture system upon which component models communicate with each other through the use of multiple, standards-based coupling interfaces whose structure and function vary in accordance with the nature of the coupling to be performed. In the most general sense, we advocate for the use of an elemental IUF (Initialize-Update-Finalize) interface for each model component but wrapped around the IUF sequence will be additional coupler functionality which mediates heterogeneity in spatial elements, time-stepping and variable semantics. Presently, there are multiple standards based frameworks that have been developed for each the elemental IUF structure and the ancillary coupler wrappers and theses are listed in Section 5 above. Effectively, the proposed solution we advocate for is not to create a single standard or framework but, instead to advocate for a coupling programming standard that works with existing and emerging coupling architectures.

The second solution we propose is to develop an Earth System Model coupling testbed to facilitate standards-based approaches to ESM component coupling and to support broader engagement and participation of geoscience domain scientists in community-based ESM development and end-use applications. This testbed will be a community facility which provides updated repositories or links to repositories of a) model component and coupling architecture code, b) code documentation, c) documentation on data semantic protocols (e.g. information on standard variable names, units and relevant metadata), and d) model benchmarking use cases from which new component model code and coupling architecture code can be evaluated. The testbed will also contain archives of users experiences in conducting model coupling activities as those histories and experiences relate to the set of ESM requirements articulated in Section 4 above, including, but not limited to the following:

- Estimates of time and resources required to complete the coupling activity (e.g. new lines of code written or person-hours committed)
- Statement of what special modifications to component model code or the coupling frameworks had to be made, if any, to complete the coupling task
- Statements regarding the basic levels of expertise or skills required
- Listings of 'roadblock' issues that arose (e.g. data congruity, platform compatibility)
- Articulation of the opportunities or potential that was 'realized' by successful completion of the coupling activity (e.g. enabling conservative, multi-scale model functionality)
- Provision of quantitative information on computational performance of systems developed
- Comments on portability of system

To initiate these solutions, this EAGER project will conduct a prototype pilot project which assesses the viability of adopting the proposed layered-architecture approach from the perspective of coupling of a set of different Earth System component models into community-based ESMs. Our solution is to conduct a pilot project which consists of a set of (to be fully defined) actual and virtual model coupling exercises that would be required to address a real-world Earth System problem (also referred to as 'use cases'). The problem and the set of models to be evaluated in this EAGER phase of the project are described in detail in the following section (Section 7. Processes).

Lastly, we suggest that to implement these solutions, NSF could play a key role in promoting interaction between the various modeling framework projects (e.g. through funding for small, focused meetings and workshops) which would help with adoption and development of common standards (e.g. for semantic mediation) as well as with community-wide buy-in to some core standards and methods.

7. ***Process:*** *Process(es) to select and evaluate standards, protocols, test data, use cases, etc. that are necessary to mature the functionality of your concept area and promote interoperability and integration between elements of EarthCube.*

In this section we articulate the details by which model coupling efforts, to be engaged in through the execution of our proposed 2-solution strategy, are to be evaluated. Recalling that the main solutions are to develop a) a multi-layer coupling architecture and b) a model coupling testbed, we require a process by which different model coupling activities can evaluated. Our goal here is not necessarily to pick 'winners' and 'losers' because any such classification could be rendered irrelevant by future development efforts. Instead, the goal of the evaluation process is to provide a reproduceable framework for harvesting and archiving the experiences of model coupling activities rather than having those experiences become lost immediately after a coupling exercise is complete, despite its failure or success. Therefore the focal point of our 'process' is to develop a simple worksheet or survey which will be completed as the coupling project proceeds. We seek to develop standardized questions to obtain input which would be utilizeable to the immediate activity as well as a broader suite of coupling activities. This worksheet will be drafted and evolved through the execution of a set of pilot studies or 'use cases' in model coupling. This section describes our progress to date in defining a relevant use case and a set of basic questions or issues to be addresses in our evaluation worksheet. While our stated use-case, to be executed during the EAGER phase of this EarthCube project, may be somewhat domain specific, focusing on water and climate interactions, our aim is the definition of a generalizable process for documenting model coupling experiences.

Use Case Description: Sustainability of Southern California Water Supply

In order to assess model coupling strategies, we propose a series of case studies focused on the sustainability of water resources in the western United States. By experimenting with and assessing various coupled models of the land surface, atmosphere, and water management systems we aim to document various advantages and disadvantages in modeling strategies used to evaluate the long term ability of California's current water infrastructure to meet the demands of agricultural and urban users, as well as the environment. California relies heavily on a system of reservoirs and aqueducts to store and transport water from the wetter, northern part of the state to the drier central and southern regions. Much of the state's water originates as snow in the Sierra Nevada Mountains, and climate change studies have predicted either significant declines or changes in the seasonality of this resource over the next century. This will undoubtedly place significant stress on the agriculture industry, which is the dominant water user in California. Additionally, the modeling systems experimented with should be able to answer questions about the short-term variability of river flow, with applications to flashflooding and hydropower. By incorporating these systems into a new, fully-coupled model framework, we can better understand the feedbacks of human water management in a region highly vulnerable to climate change

In order to capture the broad range of spatial and temporal variability required to address these problems, we will need to implement multiple suites of coupled models. The first step in this process will be to identify appropriate models and modeling systems, and their appropriate temporal and spatial scales. Here we list a set of models which have various attractive attributes or capabilities to help address our use-case problem:

1. Large-scale low-frequency hydro-limatic variability
    a. CESM – Community Earth System Model
    b. CamaFlood (global river transport/inundation model)
2. Small scale, high-frequency hydrometeorological variability
    a. WRF - Weather Research and Forecasting Model with the WRF-Hydro extension
    b. CVHM - Central Valley Hydrologic Model-- regional-scale surface and groundwater hydrology of CA Central Valley, including agriculture
    c. CHARMS - Catchment-based Hydrologic and Routing Modeling System
    d. BreZo - fine mesh, hydrodynamic flood simulation
3. Management/decision support models - Identify key components for management/decision support- what level of detail is required
    a. WEAP
    b. CalSim II - official CA DWR model for CA water management
    c. CALVIN

Coupling evaluation strategy

In addition to choosing models, we must also choose a model coupling strategy. To do this, we not only have to identify the capabilities and deficiencies of each coupling system, but also evaluate the feasibility of applying it to a set of scientific models we wish to use. This is no doubt a two-way restriction, because our choice of model-coupling strategy will in some ways constrain our choice of appropriate models, and vice-versa. Nevertheless, the principal coupling architectures (which are each defined in Section 5 above) to be explored are:

1. ESMF
2. CSDMS and vicariously, CCA
3. MCT

To evaluate the relative capabilities of these modeling systems, we define a set of metrics which allow us to compare the relative merits of each modeling system. The metrics will be based on the following criteria:

- adaptability across spatial and temporal scales

- modularity of component models
- adaptability across different computational environments including HPC environments
- integration of data from standardized data architectures such as CDF, HIS and other sources
- model performance
- scientific accuracy
- time and expertise required to conceptualize and complete the project

We are presently in the process of drafting our coupling 'worksheet' and are using the above topics as the basis for our questions. Throughout this process, we will continually assess the abilities of the chosen models to fit meet these criteria, which can broadly be grouped into 'scientific' and 'software' metrics. For example, 'modularity of component models' is a software metric, while 'adaptability across spatial and temporal scales' is a scientific metric. Balancing these sets of requirements will be key. Nevertheless, we stress again that our aim with this EAGER grant coupling exercise or 'use case' is to define a coupling evaluation and documentation process which is generalizeable enough to utilize in our proposed coupling testbed and to help guide the development of our proposed multi-layered coupling architecture. We acknowledge that this particular used case is only one example of a suite of Earth System modeling use cases that has been proposed by the broader EarthCube community such as Hypoxia of the Gulf of Mexico, Volcanic Ash Dispersal, Fire Impacts on Air Quality, Radar Impacts on Flash Flood Predictions, Forest Carbon Accounting, Arctic Coastal Processes, Hydro-Ecology and Regional Water Supply. The decision to focus on a focused issue under the general subject of regional water supply was simply a practical compromise that allowed assessment of the various Earth System Model components with which the team members have the most experience. Future efforts should attempt to apply the proposed assessment framework to these other possible use cases.

8. ***Timeline:*** *Timeline for the project and all related sub-projects, including prioritization of activities and measurable milestones/major achievements and total resources (human and financial) required to achieve roadmap goals over a period of 3 to 5 years.*

In this section we outline a basic timeline of activities that would be required to enact our proposed 2-solution strategy which is articulated in Section 6 above. The management and oversight process for our strategy is discussed in Section 9 below.

What have become Earth System models today, have evolved for over 40 years since early global circulation models were developed in the 1960's. While we don't suggest here that the challenge of Earth Systems modeling will be met within 3-5 years we do believe that truly transformative innovations in cyberinfrastructure related to Earth Systems modeling can occur during this timeframe and we propose here that a suitable coupling framework for Earth Systems modeling will emerge. Specifically, the goal of this Roadmap is identify, develop and integrate a set of standards-based protocols for linking component models of myriad Earth System processes into a coherent, conservative, multi-scale and high-performance computing enabled modeling system. Several mature and newly emerging model coupling frameworks currently exist which each may play a significant role in the future evolution of Earth System models. As such, this Roadmap advocates for the formation of a comprehensive model coupling testbed to facilitate and support coupling among different geoscience and societal component models, and to document strengths, weaknesses and opportunities with respect to the model coupling activities and the requirements of new models described in Section 4 above. Given the constant evolution of scientific understanding and the ever-increasing complexities of ESMs themselves and end-user needs we propose that such an evaluation and documentation process be considered to be evolutionary itself. Meaning, the process of benchmarking and assessing model coupling strategies for ESMs needs to evolve with the models themselves and that a process for continued support and evaluation of model coupling activities is perhaps as critical than any single model assessment effort.

This is why in addition to developing a new layered architecture for ESM coupling, which is the first proposed solution of this Roadmap, we have proposed that an evolutionary coupling testbed environment also be developed. With the active participation of key model developers and users, we expect that the development of a comprehensive and adaptable model coupling testbed process will require on the order of 3-4 years to complete. This Roadmap and the pilot project to be conducted within this existing EarthCube EAGER project will begin to lay a foundation for the articulation of such an ongoing process. We envision a project to develop an ESM-coupling testbed to have the following components and would require the associated amount of time to carry out:

1. Articulation and refinement of testbed requirements (6-9 months)
2. Compilation of leading component models and coupling architectures (6-9 months)
3. Development of quantitative metrics for system benchmarking (6 months)
4. Development of a broad class of real-world coupling exercises to serve as benchmarks (9-12 months)
5. Demonstration of a first generation coupling benchmark activities (6 months)
6. Evaluation and synthesis of benchmark activities (6 months)
7. Refinement of the overall testbed architecture (6 months)

Because some of these activities must happen in sequence, as opposed to be conducted concurrently, we expect that the entire effort would require between 3-4 years to complete. The deliverable from this initial assessment effort would be the articulation of a suite of model testbed attributes and metrics that are most suitable for Earth System Modeling applications. However, once the process is finished, a basic documentation and evaluation framework for assessing model coupling activities would also be in place and, if designed correctly, the framework will have an evolutionary capability which allows it to change with the changing requirements of ESMs.

Milestones

The milestones for tracking and documenting progress for enacting the 2-solution strategy have yet to be more fully developed but are derived from the basic timeline above. As such the specific milestones are similar to those activities specified in the timeline as follows:

1. Development of a draft publication of both coupling architecture and coupling testbed requirements and posting for community feedback
2. Establishment of component model and coupling architecture code repositories into a unified, online, searchable database
3. Development of a draft publication on quantitative metrics for system benchmarking
4. Establishment of an online, searchable archive benchmark use case descriptions and data sets which are to be utilized.
5. Acquisition and activation of a community computational facility with model and IT support to facilitate coupling testbed activities
6. Delivery and posting of revised documents on requirements and metrics

Resource Requirements

The resources required to implement and sustain the 2-solution strategy proposed here will be substantial both in logistics and in budgetary support. The challenges articulated in Section 3 above and several of the related risks listed in Section 10 below, relate to the ever increasing pace of scientific and technological innovation that is occurring with respect to a broad range of environmental models. Thus, for this particular model coupling effort to be and remain relevant significant resources will need to be brought to bear. Initially, we foresee the largest resource required is time and salary support for domain experts to develop meaningful collaborations so that they may work towards the common goals laid out by the 2-solution strategy. Thus significant amounts of support for participants' time as well as travel to collaborative meetings and workshop must be allocated. Following the drafting phase of activities, a certain level of sustained funding will be required to support the core coupling framework groups to continue to work collaboratively and also provide a broader level of support for the use of their tools. Resources for community 'support' activities are crucial to linking developers with users of various ESM components. Lastly, we note that all modeling activities hinge critically some level of computational resource. Therefore we advocate that an appropriate computational and data management architecture be acquired and maintained for the dedicated purpose of hosting the ESM coupling testbed. Without such a dedicated hardware resource and the necessary staff to support it, there is minimal chance of success in implementing the testbed and, ultimately, in providing a community pathway for more broadly engaging in model coupling activities.

9. ***Management:*** *Management/governance/coordination plan and decision-making processes necessary to successfully develop, test, and disseminate your concept. Other management considerations are to identify and respond to shifts in technologies, changing needs at the end-point of use, and integrating into the development of other aspects of the EarthCube*

cyberinfrastructure portfolio. Include discussion of approaches to educating and train end-users and evaluating community acceptance and the project responsiveness to community input.

In this section we describe what structure of project management would be required to enact the proposed 2-solution strategy. The proposed management structure is narrowly focused on elements of the 2-solution strategy and not on governing broader activities in Earth System science or related cyberinfrastructure. In particular, issues related to the governance of broader cyberinfrastrucure development for the geosciences are discussed elsewhere within EarthCube. As such the proposed management structure is intended to be succinct and focused.

1. Oversight of the layered architecture for ESM coupling and the ESM coupling testbed should be managed through the formation of a rotating steering committee comprised of developers and ‘knowledgeable’ users. The explicit role of this committee would be to review the status and evolution of the layered architecture as it evolves, to provide oversight (though not direct operational management) to the ESM coupling testbed and to liaise with funding program managers to advocate for support for various ESM coupling activities. This steering committee would lead in the drafting of requirements and metrics publications referred to in Section 8 above. It is acknowledged that mixing a broad diversity of developers and users into a single steering committee is likely to result in incoherence, since they all have separate concerns. However, in the spirit of EarthCube we feel that potential conflicts of interest and diversity of opinion can be managed through an open and transparent deliberation process.
2. Innovation of new model coupling technologies should evolve through an open, community working group or forum of participants from the entire enterprise for the purpose of establishing a participatory mechanism for model development and model use activities. This working group would be the main community avenue for participation and would report to the aforementioned steering committee on the status and emergence of coupling technologies.
3. A core set of operational staff should be directed for operational management and support of the ESM coupling testbed center. This group would lead in the archival of component code, coupling frameworks, benchmark datasets and use cases and provide overall IT support for the testbed. A subset of this operational staff in collaboration with other code developers would be responsible for developing training and support materials for model developers and end-users.

10. ***Risks:*** *Identification of risks and additional challenges to the successful implementation of the technology/approach and discussion of how those risks might be mitigated.*

There are numerous risks that could potentially impede the development of a comprehensive ESM coupling framework and the process by which such a system is utilized and evaluated. A summary of these perceived risks and potential mitigation strategies where relevant are provided in this section. As feedback on the Roadmap evolves this section will be augmented with new risks that are identified.

1. *Risk:* Science and technology will proceed at a pace that is too fast to successfully integrate into a coherent coupling framework or to reliably assess. *Mitigation strategy:* Incentivize, through funded support, model development and support centers as well as investigator teams to participate in model coupling and benchmarking activities through the proposed ESM coupling testbed.

2. *Risk:* Competition among model development groups both domestically and internationally will inhibit collaborative development, integration and assessment activities. *Mitigation strategy:* As with #1, incentivize collaboration through specific collaboration funding opportunities to encourage synthesis and synergy in parallel lines of development.

3. *Risk:* Budgetary support for integration and assessment activities is often tenuous and placed at a lower priority than foundational research or more 'edgy' model development activities. *Mitigation strategy:* Incentivize through funding opportunities for foundational model development activities to also consider and address model coupling architecture issues.

4. *Risk:* In broadening access to ESMs through multi-scale, multi-language and multi-platform functionalities, there is significant risk of ESMs being misused in a variety of manners. Modern and future ESMs are highly complex computational tools whose scientific basis and component interactions are very difficult to understand even for relevant professionals. *Mitigation strategy*: Through support of the testbed center, foster a sustained focus on development of standards for model metadata and a Community of Practice Resource maintained by the center that fully archives and documents coupling efforts. As suggested in the Sections 6 and 8 a key resource of the testbed center would be a focus on "model or coupling successes and failures" and explicit identification of characteristics or situations that have shown instability.