Need for a METOC GRID Computing Environment

James H. Corbin, Mississippi State University   Avichal Mehra, Mississippi State University
Sachin K. Bhate, Mississippi State University

Geo Resources Institute
Mississippi State University
Bldg. 1103, Rm. 233
Stennis Space Center, MS 39529
corbin@gri.msstate.edu

Abstract - **Our nation has begun major investments of time and money into an Integrated Ocean Observing System (IOOS) that will span the country's boundaries and integrate well into the international version, the Global Ocean Observing System (GOOS). Further, it has been noted that equally important is the need to integrate data management, model execution and validation, and product distribution into a modern framework or infrastructure. Every day more effort is invested in numerical ocean model R&D and then for successful R&D models there is the transition to, and implementation within, the operational activities. In a recent report sponsored by the Office of Naval Research and the National Science Foundation, it states "The field of ocean sciences is currently faced with *extraordinary opportunities for progress*... Unfortunately, many aspects of these activities are severely limited by the Information Technology Infrastructure (ITI) available to the ocean sciences community..."**

This paper discusses the need to create an infrastructure, much like the Web, that spans the needs of the ocean community from R&D, through transition and implementation, to daily operations, and that also would support distributed collaboration, education and training. This infrastructure should also provide the integrating framework for data distribution and management, model and application communication, and product dissemination and quality assurance. Common functions would be accessed the same way and have the same look and feel wherever they are encountered within the framework and thus ease the learning curve, as well as the development and collaboration process.

While this paper will suggest, at a high level, possible technology to accomplish this, the focus will be on the need for this common infrastructure and not on details of the technology to implement it. The primary technology considered is the "GRID", the computational power grid that integrates diverse high performance computing resources across a distributed environment. However, the GRID in itself can achieve even greater utility when integrated with web and distributed object/component technologies.

I. INTRODUCTION

It would be a tremendous benefit with all the modern technology in remote and in situ sensing of the natural environment, if one could count on actual observations being sufficient to provide everything needed to give clear insight into all our environmental needs. However, even if one could take enough observations over time and in space to fully meet our immediate needs, we still have need to extend our environmental assessment into a forecast. And in reality, it is not feasible nor cost effective to believe enough observations can actually be taken. Thus it is very important to incorporate numerical modeling of the natural environment, not only to produce a forecast, but also to adequately produce a realistic assessment of the dynamics in areas of sparse or no data observations. The nation is considering major investments of time and money into an Integrated Ocean Observing System (IOOS) that will span the country's boundaries and integrate well into the international version, the Global Ocean Observing System (GOOS). In their report they state a clear requirement for an integrated component for numerical modeling [1]. The report also goes on to say that equally important is the need to integrate data management, model execution and validation, and product distribution into a modern framework or infrastructure. In another recent report sponsored by the Office of Naval Research and the National Science Foundation [2], it states, "…The field of ocean sciences is currently faced with extraordinary opportunities for progress… …Unfortunately, many aspects of these activities are severely limited by the Information Technology Infrastructure (ITI) available to the ocean sciences community..".

There is an emerging need to integrate the expertise of computer scientists with Meteorology and Oceanography (METOC) modelers and applications developers to produce an ITI that can expedite these "…extraordinary opportunities for progress…" Exploiting leading edge developments in computational grid, and other leading edge software technologies, such an ITI will integrate some of their best features to provide a multi-tiered infrastructure that links a user on a desk top via a network application to model execution on high performance computational platforms that can be distributed globally. The services that can be provided by this infrastructure will be tailored to minimize the complexity of model research, development, transition, and operation or execution.

II. GRID COMPUTING ENVIRONMENT (GCE)

Presently, most numerical models are developed within the METOC R&D community as stove-pipe systems with unique requirements for control of execution, inputs and outputs. All of which are accomplished by unique execution of scripts that usually function properly only in the original research environment within which they were developed. When the model is transitioned to the operational community, the scripts have to be re-written to accommodate the new operating environment, new data inputs, and new product outputs. As a result, many man-hours are needed just to get the new model to function in the new operating environment.

A METOC GCE, on the other hand, will provide an environment that facilitates development of new models as well as their transition, implementation and operation once validated. As new models are introduced into the operational community, each one usually uses unique applications to process input and output data, such as quality control, data interpolation schemes, and graphics. Therefore an operator must know several different ways to do the same task depending on the model of interest. Within a GCE, from the user's perspective, the same functions will be done the same way through a common graphical user interface. Thus enormous time is saved in training and indoctrination of new personnel or when introducing a new model. Secondly, while most computer software developments within the research community are beginning to use industry standards, most METOC model development does not. Developing new models within this GCE infrastructure allows the model developers to exploit the latest in industry standard software engineering technology without significant changes to their underlying code.

In a nutshell, a METOC GCE improves on all aspects of model development and operation: data acquisition and assimilation; coupling; verification and validation; and, distributed processing. Its infrastructure provides improved opportunity for collaboration and partnerships, education, training, indoctrination, public awareness/service, product distribution, and distributed management of resources and projects.

Earth scientists have long yearned for affordable computing to satisfy their insatiable hunger for massive computing resources. This massive appetite for computation is a consequence of the complexity and number of computations within the numerical processes needed to simulate the natural systems and ensure accurate nowcast/forecast capabilities. Vastly different time and length scales often characterize these systems, which demands the simultaneous resolution of physical processes over a wide range of temporal and spatial scales. Such enormous computing hardware also requires support for complex software development and maintenance. Both software and hardware high performance computing (HPC) needs can only be met by pooling of limited resources for development of unified programming standards, common model interfaces, community modeling codes, and other common tools for visualization, data analysis, grid generation etc. "GRID Technologies" have been devised to enable the sharing of such resources for HPC computation, data preparation and analysis, quality assessment, and output access by diverse value-added applications. These technologies enable clusters of servers and compute servers linked via the Internet, potentially in the millions, to make processing power available on demand, rather like water or electricity, boosting computational power while cutting costs.

While the web technologies, enterprise technologies, and the majority of the distributed technologies are developed by the industry and focus on developing solutions to enable the business applications, GRID technologies were developed in academia to support effective utilization, collaboration and communication of heterogeneous high performance, geographically distributed compute resources. GRID Technologies aim to provide secure access to diverse HPC resources shared among different organizations. An HPC resource can be hardware, software, or data. Grid Technologies include resource management and scheduling, security, information services, data management and access, and problem solving environments. Currently, most of the Grid developments are based on the Globus toolkit™ [3] developed jointly by ANL and USC, including object-oriented interfaces in Java, Python and CORBA, as well as the Open Grid Services Architecture (OGSA).

III. MULTIPLE USERS AND ROLES SUPPORTED

Broadly, three different groups of users of a fully functional METOC GCE can be identified. These are researchers, operators and managers. The research support functions require the most flexibility and demand for real time system configuration. The basic tasks for a researcher include new model development, environmental simulations, validation and transition to operations. The operational user is more focused on product generation and less on processes, but needs more automation and a tailored environment for routine run streams for both operations (to run a model or to derive products from archived model runs) and transitions. The management user focus's more on monitoring status and systems (project/program and resources) management. He or she can be from research or operations, but can also represent project, program or sponsor management, or users from the public domain. These users are distinctly high-level users of the system and would use web access primarily for information only. This information could be used for generating reports, making briefs, monitoring the status of the project and/or searching for relevant documents and references.

All the roles of potential METOC GCE users can be sub-divided in two main functions: research and development (R&D); and operations (see Fig. 1), with the management role spanning both categories. The "Management" role could also be viewed as a "Decision Maker" role whose primary access to the GCE would be to access "end" products. The role of "model transition" would involve users from both these functions and hence, provides the bridge between the activities of research and operations. The connectivity shown in-between various roles in Fig. 1 represents the typical work-flow associated with development, testing, transition and operation of METOC models. For example, "applied research" which leads to development and validation of models depends on theories put forth by researchers in basic geophysical fluid dynamics and in the area of numerical methods and computational physics.

Fig. 1. Multiple roles supported by a METOC GCE

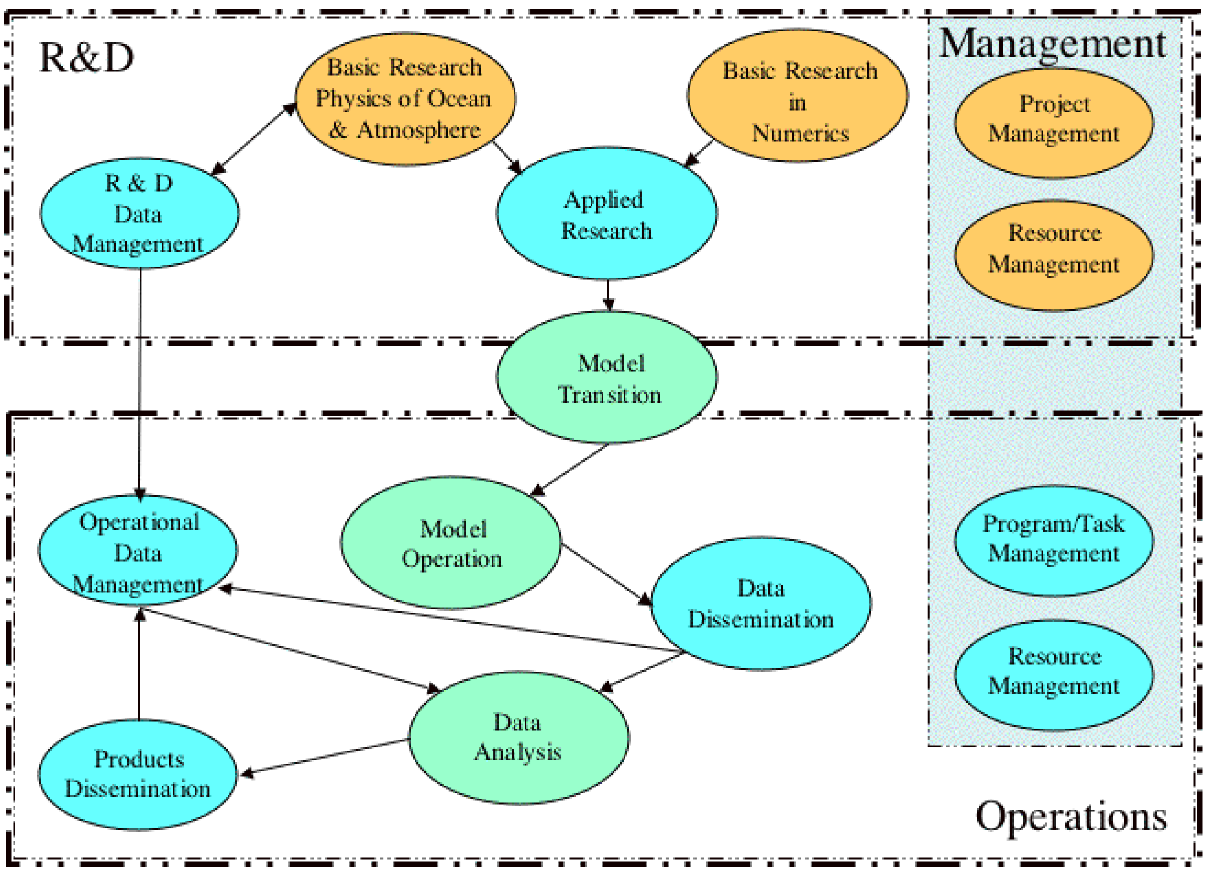

IV. METOC RELATED CYBERINFRASTRUCTURES

A "cyberinfrastructure" can be defined as an infrastructure that integrates all computing resources, including data access and management, and provides the stable, persistent, and evolutionary support required to address fundamental problems of science and technology in the present and for future.

A number of such "cyberinfrastructure" projects have been undertaken throughout the world and in this section we discuss those that relate most closely to METOC needs.

DMEFS (Distributed Marine Environment Forecast System)

The Distributed Marine-Environment Forecast System (DMEFS) infrastructure provides a distributed computing environment that focuses on METOC models and applications [4]. It provides the developer/user desktop access to geographically distributed high performance computational resources (both hardware and software), community models, distributed data servers, tools and utilities (visualization, grid generation, etc.), all in a collaborative environment via network applications. Its intelligent middle-ware assists the user (a model developer or operator) with defining a computational task, mapping the task onto computational resources, and analyzing, as well as disseminating results.

GRIP (GRid Interoperability Project)

Several different paradigms have been developed for matching resource requests with resource provision across a distributed infrastructure called a Grid. The best known of these is Globus [3] which provides functionality via a toolkit approach with API and command line interfaces. Another approach is to model the abstractions involved in the resource request and resource provision in a language such as Java so that the abstraction can be propagated to the point at which it is "incarnated", i.e. translated into the scripts necessary for it to run on a particular resource. One such system is UNICORE, developed in Europe and currently deployed across a European-wide test bed of HPC centers.

An important theoretical and practical question is whether such paradigms can interoperate, so that the user can use tools from one system to discover and use resources controlled by the other. GRid Interoperability Project (GRIP) is funded by the European Union to realize the interoperability of Globus and UNICORE and to work towards standards for interoperability in the Global Grid Forum. The GRIP project [5] will demonstrate interoperability of jobs launched via UNICORE and running on resources controlled purely by Globus. Demonstrations will be given of a portal for a relocatable weather forecasting model that can access data from large scale weather prediction models run at a large supercomputing center and tailor them to local conditions by running on local resources with results displayed to a client workstation.

PRISM (PRogram for Integrated earth System Modeling)

PRISM [6] is a pilot infrastructure project for the establishment of a climate research network. The work plan foresees the creation of a European management structure for developing, coordinating and executing a long-term program of Europe-wide, multi-institutional climate simulations. It also foresees the development of a set of portable climate community models and associated diagnostic software under standardized coding conventions that can be accessed by all European scientists and the execution of a first suite of such joint simulations. PRISM will greatly enhance the efficiency of earth system modeling in Europe and it will pave the way for the establishment of a European Climate Computing Facility. The expected product will be a flexible, efficient, portable, and user-friendly community infrastructure for Earth system modeling and climate prediction.

ESMF (Earth Science Modeling Framework)

The Earth System Modeling Framework [7] collaboration, which consists of Earth scientists and computational experts from major U.S. Earth modeling centers, is developing a robust, flexible set of software tools to enhance ease of use, performance portability, interoperability, and reuse in climate, numerical weather prediction, and data assimilation applications. The ESMF will allow diverse scientific groups to leverage common software to solve routine computational problems such as efficient data communication, model component coupling and sequencing, time management, and parameter specification. In an open dialogue with the broader community, ESMF will also develop a software interface specification so that groups working at different institutions and in different disciplines can generate interoperable software components.

IGP (Information Power Grid)

Information Power Grid [8] is the name of NASA project to build a fully distributed computing and data management environment, i.e., a Grid. The IPG project has near, medium, and long-term goals that represent a continuum of engineering, development, and research topics. The overall goal is to provide the NASA scientific and engineering communities (including the Earth Science community) a substantial increase in their ability to solve problems that depend on use of large-scale and/or dispersed resources: aggregated computing, diverse data archives, laboratory instruments and engineering test facilities, and human collaborators. The approach involves infrastructure and services that can locate, aggregate, integrate, and manage resources from across the NASA enterprise. An important aspect of IPG is to produce a common view of these resources, and at the same time provide for distributed management and local control.

V. ISSUES MINIMIZED BY GCE'S

As models and applications move through the R&D process to transition and implementation and then to daily operations, current practices require numerous modifications to code in order to accommodate differing hardware platforms, data types, data flow paths, product lines, user interfaces and requirements, and different time or scheduling constraints.

In the current way that model development is accomplished, the R&D computational environment is normally much different than the operation environment for which the model is intended. Hardware platforms may even be the same but because of differing configurations the code will need to be modified to function properly in the operational environment. The differences between the R&D environment and the operational environment are usually compensated for by the use of extensive scripting that is usually bundled outside the main model code. However, even though the scripting is done separate from the model code itself, it is usually very extensive and very time consuming to change as the model moves from R&D to operations, and even again as the operational environment changes with hardware and software updates. Further, the more complex the model or application, the more extensive and complex the scripting. And since most model configurations are distinctly different from one another, the result is the need for different people or teams of people handling each model or application.

A similar situation to the above arises with data types, data flow paths, product lines, user interfaces and requirements, and different time or scheduling constraints.

It is believed that the GCE will minimize the negative impact of the problems described above by using the middleware to provide common services for all interfacing functions and common re-useable tools and scripts that insolate the back-end changes from both the developer and the end user. And further, because the GCE provides a built-in collaborative environment, there is tremendous efficiency in providing support for education, training, indoctrination and distribution.

One last area of application that is frequently overlooked is the application of GCEs directly to the internal operating environment of the HPCs themselves. This paper has addressed the GCE as a way to insolate the user and developer from the complexities of, and changes to, the back-end processors/platforms and operating systems, but there are also tremendous economies of scale that could be realized by implementing the GCE itself within the HPC to manage hardware and software, as well as user access. There is at least one HPC already considering this kind of architecture for internal operations.

VI. CONCLUSIONS

In this paper, an attempt was not made to put an explicit "Cost vs. Benefit" analysis together regarding the above, but clearly in the final DMEFS workshop [9] it was demonstrated that a DMEFS-like infrastructure, or GCE, provides an excellent mechanism in which to integrate and share new and legacy applications. And further, that this kind of a framework provides an efficiency multiplier to all users, whether model and applications developers, operators, managers, or casual users. Models, application, data, scripts, tools, tasks, and transactions, developed or implemented by anyone within the GCE become persistent objects and can be easily shared and re-used by anyone else in the GCE by means of a simple publication process, and common functions are accessed the same way and have the same "look and feel" wherever they are encountered within the framework and thus ease the learning curve, as well as the development process. And by design, the "collaborative" nature of the GCE facilitates easy extension to education, training, indoctrination and public access. Finally, the GCE is an effective and efficient integration of the computer science IT specialist with the METOC, or other science area, specialist, scientist, and model or applications developer, and could lead to "…**extraordinary opportunities for progress…**" in many METOC fields.

References

[1] "An Integrated and Sustained Ocean Observing System (IOOS) for the United States: Design and Implementation", a report in PDF format available at http://www.ocean.us.net/projects/wspost.jsp

[2] "An Information Technology Infrastructure Plan to Advance Ocean Sciences", a report available at http://www.geo-prose.com/oiti/report.html

[3] The GLOBUS project http://www.globus.org

[4] A. Mehra, J.H. Corbin, T. Haupt, and S.K. Bhate, "DMEFS: A Web Portal for METOC Modeling", Proceedings of MTS/IEEE Oceans 2002, 29-31 October, Biloxi, MS, pp. 884-890.

[5] M. Rabbadt and P. Wieder, "UNICORE - Globus: Interoperability of Grid Infrastructures", Published in the Cray User Group Summit 2002 Proceedings.

[6] http://www.prism.enes.org/

[7] R.E. Dickenson, S.E. Zebiak, J.L. Anderson, M.L. Blackmon, C. DeLuca, T.F. Hogan, M. Iredell, M. Ji, R. Rood, M.J. Suarez and K.E. Taylor, "How Can We Advance Our Weather and Climate Models as a Community?", Bulletin of the American Meteorological Society, Volume 83, Number 3, March 2002

[8] http://www.ipg.nasa.gov/

[9] Report on Final DMEFS Workshop, available at: ftp://ftp.ssc.erc.msstate.edu/pub/mehra/DMEFS/Final_Worskhop.doc